hexagon:2611.01098v1

Fractional moments of the sharp Möbius divisor sum on integers and reflected primes

Wren Aquila

26 November 2026

Abstract

Let $M_R(n) = \sum_{d|n,\, d\le R} \mu(d)$ and $Y = \min(R, X/R)$. We determine the exponent of $\log Y$ in every absolute moment of order $0 < s \le 2$ as $Y \to \infty$. Up to the transition $s_* = 2\log(1/\log 2)/\log 2$ the exponent is the exponent for the support of M_R. Strictly below the transition the mean has the order of that support, including its factor $(\log\log Y)^{-3/2}$. Above it the exponent is $2^{s/2} - 1 - s/2$. The same laws hold on $n = N - p$ for an individual even N with fixed radical, at either $R = Y$ or $R = N/Y$ and $Y \le N^{1/2-\epsilon}$. With probability proportional to $M_R(n)^2$, the small-prime-factor count and $\log|M_R(n)|$ have the same normalized Gaussian fluctuation for $Y \le N^{1/4-\epsilon}$. The arguments use the Fourier transform of a signed divisor sum, localized divisor counts, and a lower sieve on the prime variable.

The unsigned antecedent. Let $T_y(n) = \#\{d \mid n : y < d \le 2y\}$, $\ell = \log\min(y, x/(2y)) \to \infty$, $U = \log\ell$, and $q_* = \log(1/\log 2)/\log 2$. Integrating Ford–Gabdullin–Pratt's Theorem 1.1 gives the order $\ell^{-\delta}U^{-3/2}$ for the ordinary mean of T_y^q at $0 < q < q_*$, $\ell^{-\delta}U^{-1/2}$ at equality, and ℓ^{2^q-1-q} at $q_* < q < 1$, the endpoint $q = 1$ being the elementary first moment. Thus the two branches below, and their meeting point $s_* = 2q_*$, have an unsigned precedent at half the signed moment order. In the tail integral put $r = \ell^b$. Its exponential rate is $qb - Q((1+b)/\log 2)$, where $Q(t) = t\log t - t + 1$. The maximizer leaves $b = 0$ at q_*. At that point the factor $1 + bU$ in the tail supplies the stated critical power. This same rate curve occurs in Tenenbaum's width estimate recorded in Ford (1.3). Hall–Tenenbaum had already obtained $2^q - 1 - q$ for the unsigned divisor-concentration supremum at $q \ge 1$. [3, Theorem 1.1] [6, (1.3)] [2, Theorem 3]

The signed question. For odd m, $M_R(2m) = \sum_{d|m,\ R/2<d\le R} \mu(d)$. Its value depends on cancellation among the qualifying divisors. Dress–Iwaniec–Tenenbaum established the ordinary quadratic endpoint, and Granville–Koukoulopoulos–Maynard determined support and even-moment estimates for this sharp sum. Here we bound its fractional absolute moments on integers and on $N - p$, obtain matching small moments in the ranges below and prescribe finite lists of signed values on a single reflected-prime row. Under sampling by M_R^2 the large values admit the joint limit stated below. [1] [7, (1.4), Theorems 1.3 and 1.6]

Throughout log is natural and $\log_2 t = \log t/\log 2$. Set

$$\beta = 1/\log 2, \qquad s_* = 2\log_2\beta = 1.0575\ldots, \qquad \delta = 1 - \beta + \log_2\beta = 0.08607133\ldots.$$

For $l \ge 1$ put $j_l = 1 + \log l$ and define

$$B_s(l) = \begin{cases} l^{-\delta} j_l^{-3/2}, & 0 < s < s_*, \\ l^{-\delta} j_l^{-1/2}, & s = s_*, \\ l^{2^{s/2}-1-s/2}, & s_* < s \le 2. \end{cases}$$

Fractional-moment theorem. *Let $Y = \min(R, X/R) \to \infty$ with $2 \le R \le X$. Write $\langle\cdot\rangle$ for the mean on positive integers $n \le X$. For $0 < s < s_*$,*

$$\langle |M_R|^s \rangle \asymp_s (\log Y)^{-\delta}(1+\log\log Y)^{-3/2}.$$

For $s_* \le s \le 2$,

$$\langle |M_R|^s \rangle = (\log Y)^{\,2^{s/2}-1-s/2+o(1)}.$$

Both conclusions hold with $X = N$ and the average $\pi(N)^{-1}\sum_{p<N}$ on $n = N-p$, for N even with its prime factors in a fixed finite set and $Y \le N^{1/2-\epsilon}$. They also hold for $n = p+d > 0$, $p \le X$, for a fixed nonzero even d in the same range. Constants on these rows may depend on the fixed data. At $s = s_*$ the exponent is $-\delta$; the exact power of $\log\log Y$ at this point remains open.

Theorem 1 and Corollary 8 give the small-order bounds, and Theorems 4 and 7 give the signed lower constructions. Section 11 treats $s_* \le s \le 2$.

Quadratic sampling theorem. *Fix an even squarefree D and let $N \to \infty$ through $\operatorname{rad}(N) = D$. Take $Y \to \infty$, $Y \le N^{1/4-\epsilon}$, and $R = Y$ or N/Y. On $n = N-p$ give each $p < N$ weight proportional to $M_R(n)^2$. Put $U = \log\log Y$ and $K = \#\{q \le Y : q \text{ prime},\ q \mid n\}$. Then for a standard normal Z*

$$\left(\frac{K-2U}{\sqrt{2U}}, \frac{\log|M_R| - (\log 2 - 1/2)U}{(\log 2)\sqrt{U/2}}\right) \Longrightarrow (Z, Z).$$

Both cutoffs have the same positive limiting unweighted second moment, with constant depending on D. The relation $2\log|M_R| = (\log 2)K - U + O_{\mathbb{P}}(\log(2+U))$ holds throughout $Y \le N^{1/2-\epsilon}$. The joint Gaussian limit in the interval between the quarter and half levels remains open on prime partners.

Sections 8–10 give the constant, the fluctuation law and its range. For comparison, Granville, Koukoulopoulos and Maynard give concentration of the count to relative precision on the square-weighted integers. Gorodetsky and Grimmelt obtain a Gaussian prime-factor count on reflected primes with weight $2^{\omega(N-p)}$. [7, Theorems 1.6(b) and 10.2(b)] [14]

Theorem 1. *Let $2 \le R \le X$, $Y = \min(R, X/R)$ and $l = \log(eY)$. For each fixed $0 < s \le 2$,*

$$\sum_{n \le X} |M_R(n)|^s \ll_s X B_s(l).$$

At every even integer $X = N$, writing p for primes, one has

$$\sum_{p<N} |M_R(N-p)|^s \ll_s \mathfrak{S}(N)\frac{N}{\log N}B_s(l), \quad \mathfrak{S}(N) = 2\prod_{p>2}\left(1-\frac{1}{(p-1)^2}\right)\prod_{p\mid N,\,p>2}\frac{p-1}{p-2}.$$

Both upper bounds allow a strict cutoff $d < R$.

Corollary 2. *Uniformly for $R \ge 3$ and $x \ge R^{1+\epsilon}$, and fixed $0 < s < s_*$, the mean of $|M_R|^s$ on $n \le x$ is*

$$\asymp_{s,\epsilon} (\log R)^{-\delta}(1+\log\log R)^{-3/2}.$$

In particular the conditional mean of $|M_R|^s$ on its support is bounded in this range.

The lower bound is the support lower bound immediately before (1.4) in [7], since a nonzero M_R has absolute value at least one. The authors give the support order in their displayed (1.4) for $x \geq R^{5/2}$ and obtain the lower estimate already for the range in the corollary. Hildebrand [4, section 4.3] had studied first absolute estimates for the same sharp divisor sums on rough integers. Erdős and Hall [12] showed that the density of n with $M_R(n) \neq 0$ tends to zero and bounded the real moments $\sum_{m \leq n} |M_m(n)|^q/m$, $q \geq 2$, for almost all n; Letendre [13] bounds integral moments $\int_1^n M_z(n)^t \, dz$. Apart from the Erdős–Hall density bound, these estimates average over the cutoff for a fixed integer. Our signed results average over integers for a prescribed cutoff. The lower construction in [7] uses isolated divisors in Ford's theorem [6]. Their moment theorems address even integer powers. De la Bretèche, Dress and Tenenbaum [9] give an asymptotic for the ordinary second moment, including the complementary cutoff. De la Bretèche and Tenenbaum [10] study moments of a supremum over short divisor intervals, for real orders $2t \geq 2$; their Theorem 1.5 has the lower exponent $2^t - 1 - t$ for the Möbius supremum. For $1/2 \leq t \leq 1$ their estimates give only a $(\log x)^{o(1)}$ upper bound, by Hölder from the case $t = 1$. Their remark (ii) after Theorem 1.1 makes this step for a real character, and the case $t = 1$ for μ is stated after their Theorem 1.3. The upper bound of the corollary concerns the fixed sharp cut and includes powers slightly above one. For $0 < s < 1$ it follows from the first absolute moment by integrality. We prove the upper bounds for $1 \leq s \leq 2$.

1 Lengths of translated supports

For a finite list $v = (v_1, \ldots, v_k)$ write $D_v f(x) = \sum_{A \subset [k]} (-1)^{|A|} f(x + \sum_{i \in A} v_i)$. For an interval of length H containing the support of f, the difference vanishes outside its translates by the negative subset sums. At every $w > 0$ the length of their union is at most

$$2^{\#\{i : v_i > w\}} \left(H + \sum_{v_i \leq w} v_i \right). \tag{1}$$

Divide the list at w and fix a sublist of its larger elements. This is Ford's divisor-interval bound [6, Lemma 3.1].

Consider bins $(2^{-i}, 2^{1-i}]$, $1 \leq i \leq J$, and one final bin $(0, 2^{-J}]$. Write their counts as n_i, n_0 respectively, put $S_j = \sum_{1 \leq i \leq j} n_i - j$, $m = \min_{0 \leq j \leq J} S_j$, and let t be the first minimizing index. Set $\kappa = 1 - s/2$. For $s < 2$, substituting $w = 2^{-t}$ into (1) bounds the κth power of the length by

$$2^{\kappa m} (1 + 2^J H)^\kappa \left(1 + n_0 + 2 \sum_{d=1}^{J-t} 2^{-d} n_{t+d} \right). \tag{2}$$

We used $0 < \kappa < 1$. At $s = 2$ use one for the length factor.

Lemma 3. *Sum the bound in* (2) *against the weights* $c_0^{n_0}/n_0! \prod_{i=1}^J c_i^{n_i}/n_i!$, *where* $c_0 \ll 1$ *and* $0 \leq c_i \leq c_s + e_i$, $c_s = 2^{s/2} \log 2$, $e_i \geq 0$, $c_i \ll 1$. *Apart from a factor* $(1 + 2^J H)^\kappa \exp(O_s(1 + \sum e_i))$, *the sum is at most*

$$\begin{cases} (2^J)^{\beta - q} (1 + J)^{-3/2}, & s < s_*, \ q = \log_2 \beta - s/2, \\ (2^J)^\beta (1 + J)^{-1/2}, & s = s_*, \\ (2^J)^{2^{s/2}}, & s > s_*. \end{cases}$$

Proof. The final bin sums to a constant, including its count. For the sum with a marked count n_{t+d}, fix $d \geq 1$, sum over the possible indices $t = j$ and insert one point in bin $j + d$. This costs

$c_{j+d} \ll 1$ and raises the minimum by at most one. For any initial list the earliest minimum after an insertion is either its old earliest minimum or the earliest prior index at height $m+1$. Indeed, a suffix is raised by exactly one and all earlier heights preceding the original first minimum are at least $m+1$. Thus at most two values of j qualify. Sum the geometric coefficient over d. We are left with the weight $2^{\kappa m}$. Increasing the bin rates from c_s to $c_s + e_i$ adds an independent exponential series. Inserting r points increases m by at most r, at total cost at most $\exp(2^\kappa \sum e_i)$.

If $c_s > 1$ use $m \le 0$ and sum the series to e^{Jc_s}. If $c_s \le 1$ choose q as above, so that $2^q c_s = 1$. The remaining sum equals

$$e^J 2^{-qJ} \mathbb{E}\left[2^{-(\kappa-q)(-m)-q(S_J-m)}\right]$$

with independent Poisson(1) counts. Here $\kappa - q = 1 - \log_2 \beta > 0$. For $q > 0$, fixing $-m = a$ and $S_J - m = b$ gives probability $\ll (a+1)(b+1)(1+J)^{-3/2}$. The first third of the centered walk stays above $-a$, and the increments in the reversed last third have all sums at most b. After fixing these disjoint pieces, the middle third has one specified sum, whose probability is $O((1+J)^{-1/2})$ by the Poisson mass formula. The two one-sided survival probabilities cost $O((a+1)/\sqrt{1+J})$ and $O((b+1)/\sqrt{1+J})$. Sum the exponential weights in a, b. At $q = 0$ use the one-sided bound alone for $-m$.

For that bound stop the walk or its negative on exit from $(-a-1, B)$, $B = \lceil \sqrt{J} \rceil$, assuming $a < B$. Both overshoot moments through order two are bounded uniformly conditional on either exit side. This follows from the bounded moments of the excess of a Poisson(1) variable conditional on exceeding an integer. The exit time has finite moments at fixed boundaries by runs of zero arrivals. The stopped mean martingale gives upper exit probability $O((a+1)/B)$. The stopped square minus time martingale gives mean exit time $O((a+1)B)$, by weighting the squared overshoots at both ends. Survival above $-a-1$ for J steps requires upper exit or a time greater than J, yielding the claimed estimate. □

2 Sampling the logarithms of primes

In estimates below $l = \log Y$ when Y is large. Set $v_p = \log p / l$. We use squarefree harmonic configurations of primes $p \le Y$ with one-point masses $b_p = C(p)/(p-1)$, where $C(2) = 1$ and $C(p) = (p-1)/(p-2)$ for $p > 2$. Restrictions to subsets of these primes and inflations $b_p \mapsto b_p(1+\epsilon_p)$, with $0 \le \epsilon_p \ll 1$ and $\sum b_p \epsilon_p \ll 1$, are allowed. Let $U \ge 1$ be dyadic and set $V = \min(U, l)$, $J = \lfloor \log_2 V \rfloor$. For the frequency block $U \le |\xi| \le 2U$ (including $|\xi| \le 1$ in the lowest block), the bins of Lemma 3 can be taken with

$$c_i = \left(\sum_{v_p \in I_i} b_p\right)^{1-s/2} \left(\sup_\xi \sum_{v_p \in I_i} b_p |1 - e^{i\xi v_p}|^2\right)^{s/2} \le 2^{s/2} \log 2 + e_i, \quad \sum e_i \ll 1 + \log(1 + \log^+(U/l)). \tag{3}$$

Here the prime masses may already include the allowed inflation. To prove this, put $y = l 2^{-i}$. For $y \ge 1$ the ordinary prime mass in the bin is $\log 2 + O((1+y)^{-2})$ by the prime number theorem with remainder. The adjustment from $1/p$ to $C(p)/(p-1)$ is absolutely summable. For $y \gg (1 + \log(2 + U/l))^3$, partial summation using the classical prime number theorem replaces the cosine sum by $\int_y^{2y} \cos(\xi r/l)\, dr/r$ with errors summable over the dyadic y. In detail the error is $O((1+y)^2(1+U/l)e^{-c\sqrt{y}})$ and the choice of the lower threshold absorbs $1 + U/l$ after adjusting the absolute constants. The oscillatory integral is $O(2^i/U)$ by parts, hence the total error is bounded. The other bins cost $O(1)$ each and there are $O(1 + \log(1 + \log^+(U/l)))$ of them. The inflation adds a bounded sum to the two nonnegative mass upper bounds. Expanding the two powers about the positive baselines $\log 2, 2\log 2$ gives (3).

In the final bin, for $U \le l$, use the weighted Hölder inequality with first weight Uv_p and second weight $(Uv_p)^{-(2-s)/s}$ multiplying $|1-e^{i\xi v_p}|^2$. Both prime integrals are bounded. The exponent of Uv_p in the second is $2-(2-s)/s > 0$ for $s \ge 1$, and $Uv_p \le 2$. Here $\sum_{p \le e^t} (\log p)^a/p \ll_a 1 + t^a$ for $a > 0$. For $U > l$ the final bin has $\log p \le 2$, so use both weights one. In sums of nonnegative terms the squarefree restriction only decreases the upper bound compared with factorial measures allowing repetitions.

For use with prime sampling, suppose $f(x) = B(T(x-a))$ where B runs through a fixed bounded family of compactly supported smooth functions and $T \ge 1$. The list v can have any of the above harmonic prime configurations. Let $\mathcal{C}(v)$ be the union of cells of length $1/T$ meeting the possible translated support of $D_v f$. A fixed dilation or displacement of the cell grid, or a supremum over an additional shift of size $1/T$, changes the support interval length in (1) to at most C/T. Let the sampling measure give each cell mass at most M/T. For a Fourier block F_U of $D_v f$, the integral of $|F_U|^s$ over the selected cells with respect to this measure is at most

$$C_s M |\mathcal{C}(v)|^{1-s/2} (\|F_U\|_2 + T^{-1}\|F_U'\|_2)^s. \tag{4}$$

Indeed on each cell the supremum squared is at most $O(T)$ times the integral there of $|F_U|^2 + T^{-2}|F_U'|^2$ by one-dimensional Sobolev. Use Hölder over the selected cells. Boundedly enlarged cells have bounded overlap and obey the same inequality.

Apply Hölder to the list integral at each fixed bin-count vector, with exponents $1/(1-s/2)$ and $2/s$, and use Plancherel on the second factor before taking the frequency supremum in (3). Each bin contributes $c_i^{n_i}/n_i!$. At $s = 2$ simply use the second factor. Lemma 3 bounds the total block, apart from $C_s M$, by

$$(1+V/T)^{1-s/2}(1+U/T)^s (1+\log^+(U/l))^{C_s} A_s(V) \left[\frac{\sqrt{U}}{T} (1+U/T)^{-K} \right]^s, \tag{5}$$

where K may be any fixed large integer and

$$A_s(V) = \begin{cases} V^{\beta - q}(1+\log V)^{-3/2}, & s < s_*, \\ V^{\beta}(1+\log V)^{-1/2}, & s = s_*, \\ V^{2^{s/2}}, & s > s_*. \end{cases}$$

The last factor uses the L^2 Fourier norm of f on the block. Restrict to the translated support of the full $D_v f$ before splitting it into frequencies. The sampling grids may depend on other prime configurations or translations, since (4) is uniform in them. Minkowski on the resulting joint measure sums the sth roots of (5). For $U \le T$ the growth from $\sqrt{U}$ controls the dyadic sum, and for $U > T$ choose K larger than all fixed power exponents. Thus the total is

$$\ll_s M T^{-s/2} A_s(\min(T,l))(1+\log^+(T/l))^{C_s}. \tag{6}$$

All the estimates are uniform over the admitted subsets of primes and inflations of bounded total mass.

3 Two prime scales on the integer rows

We give the details for both types of row. Put $L = \log X$, use $X = N$ on a reflected prime row, and write $W = X$ for its mass on all integers, $W = \mathfrak{S}(N)N/L$ for the normalization on reflected

primes. Thus a sum over the row of an integrable fixed smooth divisor profile at scale l is of order at most W/l. We detail the uniformity as Y varies.

We use the standard upper bounds for nonnegative multiplicative F with $F(1) = 1$, $0 \le F(p^k) \le A$:

$$\sum_{\text{row}} F(n) \ll_A \frac{W}{L} \exp\left(\sum_{p \le X} \frac{F(p)}{p}\right). \tag{7}$$

For integers this is the multiplicative mean-value upper bound. For the prime row this is the specialization of Henriot's corrected Theorem 5 [11] to $n, N-n$, with the prime slot upper-sieved. To verify the local factors, majorize the prime slot outside $p \le X^{1/10}$ by the multiplicative indicator $h(n) = 1_{P^-(n) > X^{1/10}}$. Use the primitive linear forms $u - X, 2X - u$ on $X < u \le 2X$. They have joint degree two and coefficient norm $O(X^2)$, so fixed Henriot parameters $\alpha = 1/2$, $\delta = 1/4$ suffice. Uniform class membership follows from $A^{\omega(n)} \ll_{A,e} n^e$ at each fixed $e > 0$. For a prime r, write $e_r = (1 - 1/r) \sum_{j \ge 1} F(r^j)/r^j = F(r)/r + O_A(r^{-2})$. The extended local-density sum has factor $1 + e_r + h(r)/r$ at $r \nmid X$ since the two roots are distinct; at $r \mid X$ its factor is $1 + h(r)e_r$ because both positive valuations occur together and h is strongly multiplicative. Thus its Euler product is bounded by the exponential in (7). The sieve product, with two roots off X and one on X, is $O(\mathfrak{S}(X)/L^2)$ by Mertens, absorbing the prime 2. The few small primes cost $O_A(X^{1/5})$ using the elementary growth bound. In particular $F(n) = B^{\omega(n; \le H)}$, with $H = Y^{5/4}$, has row mean $\ll_B W l^{B-1}$ for $Y \ge 3$. Here $\omega(n; \le H)$ counts distinct prime factors up to H.

For $n > 1$ write $\Delta = \log(X/\operatorname{rad} n)$. The squarefree divisor complement exchanges the closed cutoff at R with the strict cutoff at $\operatorname{rad}(n)/R$, up to the unit sign $-(-1)^{\omega(n)}$. It suffices to handle the two cut variables $\log d/l + j\Delta/l$ against the step at one, $j = 0, 1$, with $Y \le \sqrt{X}$. All effective divisors in a profile supported on the right below $5/4$ have $d \le H$. For bounded Y, the step contains only boundedly many such divisors and the assertions follow directly. The term $n = 1$ can always be treated separately.

Resolve the step on its left into one fixed smooth remainder and dyadic compact smooth bumps f of height $O(1)$ and width $O(1/T)$ around one, supported in $[3/4, 5/4]$. For instance subtract successive smooth cutoffs which make the transition to zero on $[1 - C/T, 1]$. Stop at a dyadic $T \asymp l^{50}$ and later pay for the residual window. A fixed extension to the left is harmless. For the smooth remainder, Fourier inversion gives the product $\prod_{p \mid n,\, p \le H} |1 - e^{i\xi v_p}|$ with an extra phase of modulus one at the moving cut. For $s \ge 1$ use Minkowski in the row and apply (7) to its sth power. By splitting at $v_p = 1/(1 + |\xi|)$ the exponential prime sum for $p \le H$ costs at most a fixed power of $2 + |\xi|$. Primes above H have weight one, giving $O(L/l)$. Rapid Fourier decay gives a total sth moment $O_s(W/l)$.

For a dyadic bump fix a sufficiently large absolute K_1 and successively try $u = u_0, u_0/2, \dots$ with $u_0 = 1/64$, stopping at a last value $u_{\min} \asymp 1/(K_1 \log(2l))$. At each scale put $z = Y^u$, $Z = X^u$, and require that the products of the distinct prime factors of n at most z, Z be at most $Y^{1/8}, X^{1/8}$ respectively. Use the first passing scale. If neither condition can be satisfied together at the last scale, the failing product gives a Rankin factor $e^{-1/(8u_{\min})}$ times the corresponding multiplicative inflation $p^{1/\log z}$ or $p^{1/\log Z}$ at the primes below that threshold. Its local values are at most e and $\sum_{p \le P}(p^{1/\log P} - 1)/p = O(1)$. Multiplying by the crude bound $O_s(2^{s\omega(n; \le H)})$ and using (7) gives $O(W l^{-30})$ after choosing K_1. If a scale $u < u_0$ first passes, one of the conditions failed at $2u$. The same Rankin factor is then $e^{-1/(16u)}$ with inflation below Y^{2u} or X^{2u}. Keep it in the sum. It will pay for all fixed inverse powers of u.

We can discard $n < X/Y^{1/16}$, $n/\operatorname{rad} n > Y^{1/16}$, and any repeated prime above $Z_{\min} = X^{u_{\min}}$, at a cost $O_s(W l^{-30})$ per bump. On the prime row discard also $p \le \sqrt{X}$ and $p \mid N$, where $n = N - p$. Here are the estimates. The short reflected interval costs $O(X/(LY^{1/16}))$ in prime

count by the interval sieve. The excessive square part gives a square divisor $d^2 \mid n$ with $d > Y^{1/32}$ by taking the largest square divisor. For $d \le X^{1/4}$ a reduced progression costs $O(X/(Ld\varphi(d)))$ by Brun–Titchmarsh; for $d > X^{1/4}$ sum the integer bound X/d^2. Nonunits were discarded. Also $\sum_{d>D} 1/(d\varphi(d)) \ll 1/D$ by partial summation. For a repeated prime use $D = Z_{\min}$. On integers use X/d^2 throughout. Cauchy with the $2s$th moment of $2^{\omega(n;\le H)}$ just proved makes the exponential or power savings dominate every power of l. The retained rows have $0 \le \Delta \le l/8$.

Split the divisors first at z, retaining the small-prime difference $K_m = D_{(\log \ell/l:\, \ell \mid m,\, \ell \le z)} f$. For each contributing large sublist e of primes in $(z, H]$, some squarefree divisor d giving its translate satisfies $Y^{5/8} \le d \le Y^{5/4}$. Hence e contains a prime $q > z$, and its squarefree complement $\mathrm{rad}(n)/d \ge X/Y^{11/8} \ge X^{5/16}$ contains a different prime $r > Z$. In the last assertion use the bound on the distinct product up to Z. Since $r > Z \ge z$, it is outside e. To raise the original large-sublist sum to power s, use

$$\left|\sum_e \mu(e)K(e)\right|^s \le 2^{(s-1)\omega(n;(z,H])} \sum_e |K(e)|^s.$$

In each nonzero term overcount by all admissible choices of q, r. Put $n = mqr$ and $e = qe'$. The small list is unchanged, $r \nmid mq$, and e' uses only primes of m in $(z, H]$. In the Rankin inflation from a failed scale, removing q, r loses a factor at most e^2. The extra count power changes by a bounded factor.

For clarity, fix m, q, write $D = X/(mq)$ and $h = \min(1, l/T)$. Bin the possible r by $bh \le \log(D/r) < (b+1)h$ for integers $b \ge 0$. In every nonempty bin the upper endpoint De^{-bh} is at least Z and the full bin length is $\asymp hDe^{-bh} \gg Z^{1/2}$. On a reflected prime row, sieve the two linear forms $r, N - mqr$ in this interval to a fixed small power of Z. The dimension-two upper-bound sieve gives at most

$$\ll \mathfrak{S}(N)C(m)hDe^{-bh}/(uL)^2, \qquad C(m) = \prod_{\ell \mid m,\, \ell > 2} (\ell - 1)/(\ell - 2). \tag{8}$$

Indeed the forms have two roots at each odd prime outside Nmq and one root at primes dividing Nmq, since $(mq, N) = 1$. The extra $C(q)$ is bounded. The standard Selberg upper sieve with this local product applies uniformly on any such full interval: its integer divisor errors are bounded by the number of roots and the level can be a fixed power of the interval length. All counted primes exceed the sifting limit. On the integer row sieve only r, giving $O(hDe^{-bh}/(uL))$. Keeping $C(m) \ge 1$ is harmless there.

The translated argument of K_m is

$$(1 + j\mathbf{1}_{q \mid m})\frac{\log q}{l} + \frac{\log e'}{l} + \frac{j}{l}\left(\log(m/\operatorname{rad} m) + \log(D/r)\right).$$

This accounts for a repeated q. Replacing $\log(D/r)$ by bh costs a supremum over shifts of at most $1/T$. For each bin and m, e' handle the two possible coefficients of $\log q/l$ separately and enlarge both sums to all $z < q \le H$. After (8) the sampling measure is $1/q$. It has mass $\ll 1/(uT)$ on any cell of width $1/T$ in these logarithmic coordinates, also after a dilation by two. For physical logarithmic length $l/T \le 1$ the prime interval at height $P \ge z$ has length $\asymp Pl/T$ with logarithm comparable to $\log P$, since $z = Y^u$ exceeds every fixed power of l. Brun–Titchmarsh and division by P give the bound. Clips at the endpoints may be bounded by the full cell, or by its extension from z. For longer cells partition into pieces of physical logarithmic length comparable to one and use the same bound. The cost of summing he^{-bh} is $O(1)$.

Large primes in $(z, H]$ in the harmonic sum over m carry at most 2^s per prime, including the choices of e'. Those in $(H, X]$ carry one. Thus, including $C(m)$ and either retained Rankin inflation, their Euler product is $O_s(u^{-C_s}L/l)$. Indeed the two harmonic ranges cost $2^s \log(1/u) + O_s(1)$ and

$\log(L/l)+O(1)$ in the exponent, and the inflation adds a bounded sum. The remaining small-prime sums have exactly the permitted masses $C(p)/(p-1)$ after summing arbitrary positive powers. The sampling bound is uniform in translations, even when the translation depends on the entire small list. This last uniformity follows by enclosing the selected cells using (1) with a fixed enlargement of H there; the Sobolev factor uses only the full-line norm.

Combining (6) with these factors gives, on the scale u,

$$\ll_s \frac{W}{l}\, u^{-C_s}\, T^{-s/2} A_s(\min(T,l))(1+\log^+(T/l))^{C_s}$$

for the sth moment of the bump. After the first scale, each term has the additional factor $O(e^{-1/(16u)})$. Their dyadic sum is bounded by an absolute constant depending on s. This summation is needed to preserve the negative power of $\log l$ in the answer.

4 The sharp cutoff and its endpoint

Apply Minkowski over the dyadic bumps and take their sth roots. For $T \le l$, the power of T in $T^{-s/2}A_s(T)$ is positive, and for $T > l$ there is geometric decay times a fixed power of $1+\log(T/l)$. The total sth moment is therefore

$$\ll_s \frac{W}{l}\, l^{-s/2} A_s(l) + O_s(W/l) = O_s(W B_s(l)).$$

The discarded $O(Wl^{-30})$ terms summed over $O(\log l)$ scales also fit. For the equality of powers below s_* use $\beta - q - s/2 = 1-\delta$.

It remains to justify the dyadic stopping, including the atom at the cutoff. Write $\eta = Cl^{-49}$. A nonzero residual requires an effective divisor with logarithm within η of the physical cutoff. For $j=0$ sum the individual divisor counts over $[Ye^{-\eta}, Ye^{\eta}]$. On reflected primes at unit modulus D the count is $O(X/(L\varphi(D)))$ by Brun–Titchmarsh, as $Y \le \sqrt{X}$. The sum is $O(W(\eta+1/Y)(1+\log l))$, using $D/\varphi(D) \ll 1+\log\log(3D)$. For $j=1$ take instead the complementary divisor $D = \mathrm{rad}(n)/d$ in $[(X/Y)e^{-\eta}, (X/Y)e^{\eta}]$. Write $n = kD$, with $k \le 2Y$. For each k sieve $N - kD$ on the full interval of length $\asymp \eta X/Y$, whose logarithm is $\gg L$. The cost is $O((\eta X/Y)(k/\varphi(k))/L)$. Sum using $\sum_{k\le 2Y} k/\varphi(k) \ll Y$. In both counts $p \mid N$ may be discarded at negligible cost, and in the second count so may $p \le \sqrt{X}$. On ordinary integers the corresponding bounds are $O(X(\eta+1/Y))$ and $O(\eta X + Y)$ by counting multiples. These are bounds for the number of nonzero residuals. Their size is $O(2^{\omega(n;\le H)})$. Cauchy and (7) at $2s \le 4$ give $O_s(Wl^{-15})$ for their sth moment, since the divisor moment costs at most $O(Wl^{15})$. This covers strict and closed cutoffs and proves Theorem 1.

5 Matching lower bounds on individual reflected prime rows

Theorem 4. *Fix $0 < \epsilon < 1/2$ and an even integer $D \ge 2$. Let N be even with $\mathrm{rad}\, N \mid D$. Uniformly as $Y \to \infty$, $Y \le N^{1/2-\epsilon}$, let $H_t(N,Y)$ count primes p with squarefree $n = N-p$ in $(N/2, 3N/4)$, $M_Y(n) = t$ and $|M_{N/Y}(n)| = 1$. Then for each $t = \pm 1$,*

$$H_t(N,Y) \gg_{D,\epsilon} \frac{N}{\log N} (\log Y)^{-\delta} (1+\log\log Y)^{-3/2}.$$

Consequently the orders in Corollary 2 hold on each such prime row at the two cutoffs $R = Y, N/Y$, with its mean normalized by $N/\log N$. This includes $N = 2^m$.

Ford also counts rough integers with a divisor in a fixed-ratio interval [5]. Ford [6, Theorems 6–7] and Koukoulopoulos [8] treat the existence of a divisor of a shifted prime in an interval. Here one prescribes the signed sum and uses a varying reflected total. The following construction uses the isolated divisors in Ford's (5.5).

Proof. Put $l = \log Y$ and $L = \log N$. Fix a prime $b_0 \geq 3$ coprime to D. Choose fixed $\kappa > 0$ small in terms of ϵ, then $\eta > 0$ sufficiently small in terms of D, ϵ, κ, with $\eta < \kappa/4$. Set $z = Y^{\kappa}$. A divisor $b \mid a$ is σ-isolated if all other divisors b' have $|\log(b'/b)| > \sigma$. Take a fixed $\sigma \geq 100$ with $e^{\sigma} > 10 b_0 D$. For each of the two parities of $\omega(b)$, Ford gives

$$\sum_{a \leq Y^{\eta}}^{*} \frac{1}{a} \#\{b \mid a : b \text{ is } \sigma\text{-isolated, of the prescribed parity}\} \gg_{D,\eta} l^{2-\delta}(1+\log l)^{-3/2}. \tag{9}$$

The star takes squarefree a with $P^{-}(a) > e^{\sigma}$ and $\omega(a)$ odd. To see the parity detail, use [6, p. 395, proof of (5.5)] with $g = 1$, fixed σ and $\alpha = \eta$. It supplies a single $k = v$ with $v = \lfloor \log_2 \log y - c_{\sigma,\eta} \rfloor$ in its notation (Lemma 4.8). Take $y = Y$ or $Y^{1/2}$ to make v odd, retaining the same order. Complementation $b \mapsto a/b$ preserves isolation and reverses the parity for odd v. It gives exactly half of the isolated divisors of either parity.

For each selected pair choose primes

$$\frac{Y}{b_0 b} < q < \frac{Y}{2b}$$

and put $e = b_0 a q$. The interval has fixed ratio $b_0/2 > 1$ and its left end exceeds $Y^{1-2\eta}$. By the prime number theorem its reciprocal prime sum is $\gg_{b_0} 1/l$. For one a the intervals are disjoint, by isolation. The largest prime q determines the rest of e uniquely and $(e, N) = 1$. Thus for each prescribed parity the family $\mathcal{E}$ has

$$\sum_{e \in \mathcal{E}} \frac{1}{\varphi(e)} \gg_{D,\eta} l^{1-\delta}(1+\log l)^{-3/2}, \qquad e \leq b_0 Y^{1+\eta}. \tag{10}$$

Count primes p in the displayed central interval with $N - p = ew$ and $P^{-}(w) > z$. For squarefree d formed of the sifting primes $\ell \leq z$ outside N, the logarithmically weighted prime count has main term $N/(4\varphi(ed))$. Thus the local density relative to the $d = 1$ term is $g_e(\ell) = 1/(\ell - 1)$ if $\ell \nmid e$, $1/\ell$ if $\ell \mid e$, and zero at $\ell \mid N$. A prime divisor of N cannot divide $N - p$ here for sufficiently large N, since its prime factors divide D. The dimension-one product condition is uniform in e by comparison with the product over all odd primes of $(1 - 1/(\ell - 1))^{-1}$. The lower linear sieve at level z^4 has a uniformly positive main factor for large z, and $\prod_{\ell \leq z}(1 - g_e(\ell)) \gg 1/\log z$. In the sum over e the accumulated remainder is smaller than N/L^A for any fixed A. Indeed $ed \leq b_0 Y^{1+\eta+4\kappa} < N^{1/2-c_\epsilon}$ by choice of κ, and a modulus $m = ed$ has at most $\tau(m)$ such decompositions. If E_m is the maximum absolute discrepancy on the two endpoints and reduced classes, Bombieri–Vinogradov bounds $\sum_m E_m$ with arbitrary logarithmic precision. Brun–Titchmarsh gives $E_m \ll_\epsilon N/\varphi(m)$; the Euler product gives $\sum_m \tau(m)^2 E_m \ll N L^{O(1)}$. Cauchy's inequality absorbs the divisor weight. These estimates use the prime logarithm in place of the Mangoldt function, which has the same Bombieri–Vinogradov bound (or sum the elementary prime-power error at this shortened level). Removing the logarithm on the central interval, the retained count with its current multiplicities is at least

$$\frac{cN}{L \log z} \sum_{e \in \mathcal{E}} \frac{1}{\varphi(e)}. \tag{11}$$

The constant can be fixed independently of η once κ is fixed.

Discard any w with a divisor v in $[Y^{1-\eta}/(2b_0), Y]$. Such v is z-rough. For fixed e, v, sieve the integer $u = w/v$ on the central interval of length $N/(4ev)$. Require u to have no prime $\le z$ and $N - evu$ to be prime. We have $ev \le b_0 Y^{2+\eta} \le N^{1-c_\epsilon}$. The upper-bound sieve, sifting both forms to z and the prime form to a fixed small power $N^{c'_\epsilon} \ge z$, bounds their count by

$$\ll_{D,\epsilon} \frac{N}{L \log z} \frac{1}{\varphi(e)} \frac{1}{v}.$$

For example, sieve the product with both roots up to z and the single root thereafter. Take c'_ϵ small enough for the dimension-two upper sieve with integer progression errors $O(2^{\omega(d)})$ and then shrink κ if necessary. The primes dividing Nev can save at most the bounded D-factor and $\prod_{p|ev,\ p>2}(p-1)/(p-2) \ll (e/\varphi(e))(v/\varphi(v))$; $v/\varphi(v) \ll 1$ since $v \le Y$ is Y^κ-rough. Nonunits $(v, N) > 1$ do not occur. A one-dimensional upper sieve on each dyadic interval in the indicated range gives

$$\sum_{Y^{1-\eta}/(2b_0) \le v \le Y,\ P^-(v) > z} \frac{1}{v} \ll \frac{\eta l + O_{b_0}(1)}{\log z}.$$

As the ratio is $O(\eta/\kappa) + o(1)$, this removes less than half of (11) after fixing η sufficiently small.

In any remaining representation, $b_0 a$ is exactly the part supported at primes at most z. The marked prime q is the only prime divisor of n in its possible range $(Y^{1-\eta}/b_0, Y]$, since a second one would divide w in the discarded interval. This also excludes $q \mid w$. Thus the surviving count has multiplicity one. The count of nonsquarefree partners now costs $O_D(N/(Lz) + N^{3/4})$: a repeated prime must exceed z, and one may sum Brun–Titchmarsh over its square up to $N^{1/2}$, and the elementary integer bound on multiples above that. This is negligible relative to (11) and (10) as $Y \to \infty$.

For $n = b_0 aqw$ as retained and $Y/2 \le T \le Y$, cancellation at b_0 gives

$$M_T(n) = \sum_{d \mid aqw,\ T/b_0 < d \le T} \mu(d) = \mu(qb).$$

Indeed qb is strictly inside every indicated interval. Among divisors involving q it is the only one: another factor from w is at least z, giving $qz > Y$, and the other divisors of a are excluded by isolation. A divisor avoiding q must use a nontrivial divisor $v \mid w$ of size $\ge Y/(2b_0 a)$ and $\le Y$, which was excluded. Since b has prescribed parity, $\mu(qb) = t$ is arbitrary. The endpoints have strict margins by construction, so the same identity applies to the open cutoff. Finally $nY/N \in (Y/2, 3Y/4)$ and the squarefree complement gives $M_{N/Y}(n) = -\mu(n)\, M_{<nY/N}(n)$. Equations (11) and (10) finish. □

6 Local factors and totals divisible by the first primes

Write $P(w) = \prod_{p \le w} p$ and $B(l) = l^{-\delta}(1 + \log l)^{-3/2}$. In this section $l = \log Y$, $L = \log N$, and all estimates with H_t have $Y \le N^{1/2-\epsilon}$ for fixed $\epsilon > 0$.

Theorem 5. *Suppose $P(w) \mid N$, $w \ge 3$ and $u = l/\log w \ge 2$. For either $R = Y, N/Y$ and $0 < s \le 2$ one has*

$$\sum_{p<N} |M_R(N-p)|^s \ll_s \frac{\mathfrak{S}(N) N}{L \log w} B_s(u).$$

If $\operatorname{rad} N = P(w)$ *and* $u \to \infty$*, for each* $t = \pm 1$*,*

$$H_t(N, Y) \gg_\epsilon \frac{\mathfrak{S}(N)N}{L \log w}\, B(u) \asymp \frac{N}{L} B(u).$$

For arbitrary even N*, if* $l/\log L \to \infty$*, then*

$$H_t(N, Y) \gg_\epsilon \frac{\mathfrak{S}(N)N}{L \log L}\, B\Big(\frac{l}{\log L}\Big).$$

In particular both signs occur at every such even total.

Proof. In the upper bound $p \mid N$ contributes at most $N^{o(1)}$ by the elementary divisor bound. On all remaining rows the small prime list has no member at most w. Use $V = \min(U, l/\log w)$ in the bin calculation. When $U > l/\log w$ the final bin contains only primes in $(w, w^2]$, of bounded harmonic mass. When $U \le l/\log w$ keep the earlier weighted estimate. The oscillatory errors on the other bins are bounded by the earlier error sum, so (6) holds with $l/\log w$ as its frequency cap. Outside the small list we used $O(u_0^{-C_s} L/l)$ for the large-prime factor at a trial scale u_0, which remains valid by enlargement. Summing the bumps gives $(\mathfrak{S}(N)N/(Ll))(l/\log w)B_s(l/\log w)$. The previous remainders $O_s(\mathfrak{S}(N)N/(Ll))$ fit since $vB_s(v) \gg_s 1$ for $v \ge 2$.

For the lower bounds use $W = \max(20, w)$ in the first case, $W = L^2$ in the unrestricted case. Then $\log W = o(l)$. In the earlier construction allow each prime $b_0 \in (W, W^2]$ with $b_0 \nmid N$, and take $\sigma = \log(11W^2)$. Use the uniform version of Ford (5.5), with its parameter $\sigma = \eta_{\mathrm{Ford}}$ and $u_{\mathrm{Ford}} = \sigma/\log y$. The implied constant there depends only on our fixed exponent η in $a \le Y^\eta$. Ford states (5.5) for $\eta_{\mathrm{Ford}} \ge 100$; when $\sigma < 100$ apply it at $\eta_{\mathrm{Ford}} = 100$ instead, since a divisor isolated at width 100 is isolated at width σ and $P^-(a) > e^{100}$ implies $P^-(a) > e^\sigma$. The same two choices $y = Y, Y^{1/2}$ make the selected number of primes odd. For either prescribed divisor parity, the isolated harmonic sum is

$$\gg_\eta (l/\log W)^{2-\delta}(1 + \log(l/\log W))^{-3/2}.$$

It already has $(a, N) = 1$ in the primorial case. In the unrestricted case remove a sharing a prime $r > 11W^2$ with N. The cost is $\ll l^2 \sum_{r \mid N, r > 11W^2} 1/r = o(1)$: on writing $a = ra'$, use $I(a;\sigma) \le \tau(a) \le 2\tau(a')$ and the harmonic divisor bound. The number of possible r is at most $L/\log 2$.

For each such a, b, b_0, the reciprocal prime mass in $(Y/(b_0 b), Y/(2b))$ is $\gg \log b_0/l$, by the ordinary prime number theorem on multiplicative intervals of fixed ratio and summation. In the unrestricted case the primes dividing N cost $O(LW^2/Y^{1-\eta}) = o(1/l)$. In the primorial case all these primes are already larger than w. Also $\sum 1/b_0 \gg 1$ on the allowed interval: in the unrestricted case the excluded reciprocal mass is $O(L/W)$. Thus (10) now has the right side

$$\gg_\eta (l/\log W)^{1-\delta}(1 + \log(l/\log W))^{-3/2}.$$

The smallest factor b_0 is unique. The progression levels have $e \le W^2 Y^{1+\eta} \le Y^{1+2\eta}$. Replace the lower endpoint of the discarded v-interval by $Y^{1-\eta}/(2W^2)$. Its total harmonic cost is $O((\eta + O(\log W/l))/\kappa)$. The small part and q are still unique; in particular $qz > Y$.

The lower sieve product now retains a factor $\gg \mathfrak{S}(N)$ from the primes dividing N up to z. In the primorial case $w < z$, and in the unrestricted case $\sum_{r \mid N,\, r > z} 1/r \le L/z = o(1)$. The upper sieve on the discarded v has the same factor $O(\mathfrak{S}(N))$ with an absolute implied constant, by its local roots. Primes $p \mid N$ in the original central interval may be omitted at each stage. For $(ed, N) = 1$

their removal in Bombieri–Vinogradov costs $O(L^2)$ per modulus, and on the remaining rows every divisor of $N-p$ is a unit modulo the prime factors of N. All constants in the sieve comparison and hence the choice of η are uniform. The squarefree exceptions are negligible as before. Multiplying the new mass by $\mathfrak{S}(N)N/(Ll)$ proves the assertion. For $\operatorname{rad} N = P(w)$, Mertens' theorem gives $\mathfrak{S}(N) \asymp \log w$. □

Corollary 6. *If $\mathfrak{S}(N) \le C$ and $(\log N)^{1+\xi} \le Y \le N^{1/2-\epsilon}$ with fixed $C, \xi, \epsilon > 0$, then $H_t(N,Y) \gg_{C,\xi,\epsilon} (N/L)B(l)$ for both signs. The same bound holds when $\omega(N) \le K$ for fixed K, throughout $Y \to \infty$, $Y \le N^{1/2-\epsilon}$.*

Proof. Choose primes $b_0 \ge 3$ and $r_0 > 4b_0$, both outside N and bounded in terms of C. This is possible because $\sum_{r|N} 1/r$ is bounded in terms of C, and the sum of reciprocals of primes diverges. Choose $\sigma \ge 100$ fixed so $e^\sigma > 10b_0r_0$. Applying Ford with $a \le Y^{\eta/2}$, discard prime factors of N from each a individually. An isolated divisor maps to an isolated divisor of $a/(a,N)$ on dividing out the same primes. For a deleted part $d \mid \operatorname{rad} N$ a resulting divisor has at most $2^{\omega(d)}$ preimages. The unrestricted harmonic sum is therefore at most $\prod_{r|N}(1+2/r) \ll_C 1$ times the isolated harmonic sum on a coprime to N. No parity needs to survive this deletion. Adjoin r_0 to each resulting a. For each formerly isolated b, both b and r_0b are isolated at width $\log(3b_0)$ in r_0a and have opposite parity. This costs only the fixed factor $1/r_0$. It yields (9) for each desired sign. The new a is at most Y^η, and it remains coprime to b_0N. Primes $q \mid N$ in a marked interval cost $O(L/Y^{1-\eta}) = o(1/l)$ on shrinking η in terms of ξ. Under $\omega(N) \le K$ their cost is $O(K/Y^{1-\eta})$ instead. The remaining estimates in Theorem 4 are uniform since $\mathfrak{S}(N) \le C$; for the lower sieve it suffices to retain its absolute product lower bound. This proves the corollary. □

7 Prescribing finitely many values on the same partner

Theorem 7. *Fix distinct positive real numbers $t_1, \dots, t_h$, two vectors $k, m \in \mathbb{Z}^h$, an even integer D and $\epsilon > 0$. Let N be even with $\operatorname{rad} N \mid D$, $Y \to \infty$, $Y \le N^{1/2-\epsilon}$. At least*

$$\gg_{D,\epsilon,t,k,m} (N/\log N)B(\log Y)$$

primes p have a squarefree partner $n = N - p \in (N/2, 3N/4)$ satisfying simultaneously, for all i,

$$M_{t_iY}(n) = k_i, \qquad M_{N/(t_iY)}(n) = -\mu(n)m_i.$$

If any prescribed entry is nonzero, this count also has an upper bound of the same order. The assertions hold on the fixed shifted primes $n = p + d > 0$ for nonzero even d, using X as the upper total and $D = \operatorname{rad}|d|$.

Proof. Choose $1/2 < \lambda < 3/4$ so that all $2h$ positive numbers $t_i, \lambda t_i$ are distinct. We first prescribe the values k_i, m_i at these small cutoffs, throughout a fixed relative neighborhood of each cutoff. Take the primes with n/N in a sufficiently small fixed interval about λ. Squarefree complementation then gives the displayed simultaneous values.

Order the $2h$ desired locations as $0 < x_1 < \cdots < x_j$ and write their prescribed integer values as $c_1, \dots, c_j$, with $c_0 = 0$. Fix a prime $b_0 > 10x_j/x_1$ outside D. In each consecutive gap choose y_i strictly between x_{i-1} and x_i, with $y_1 > x_1/2$ for the first gap. Choose a second prime $r > 10b_0x_j/y_1$ outside b_0D; increase it further if needed so $x_j/r < x_1/(3b_0)$. Choose a sufficiently large fixed real A. At each y_iA install $|c_i - c_{i-1}|$ squarefree divisors with the required sign as follows. For a negative increment take distinct primes in a very short fixed relative interval around y_iA. For a positive

increment take distinct primes near y_iA/r, and multiply each by the common r. All primes can avoid b_0rD by the prime number theorem and all can be chosen distinct. Let f be r times the product of the chosen primes. These designated divisors are the only divisors of f in $[x_1A/(2b_0), 2x_jA]$, once $A > \max(3x_jr^2/y_1^2, 2b_0r/x_1)$ with r already fixed and the prime intervals of sufficiently small fixed relative width: the single primes for positive increments lie below this interval, r times the single primes for negative increments lie above it, and any two chosen primes have product $\gg A^2/r^2 > 2x_jA$. The prime r by itself and 1 are below the interval. In particular all lower endpoints near x_iA/b_0 lie below every designated divisor. The widths of the prime intervals and of all cutoff neighborhoods may be chosen fixed and small enough to preserve the strict inequalities. Zero increments require no designated prime.

Take a squarefree $a \le Y^\eta$ and an isolated $b \mid a$ as before with $\omega(b)$ odd, so that $\mu(qb) = 1$ for a new prime q. Use a fixed isolation width so large that $P^-(a) > b_0fD$ and that the entire interval $[x_1A/(3b_0f), 3x_jA]$ lies inside $(e^{-\sigma}, e^\sigma)$. For each a, b take primes in a short fixed relative interval about $q = Y/(Ab)$. This keeps all qb in the prescribed neighborhoods after rescaling. The intervals are disjoint for one a; their reciprocal prime mass is $\gg 1/\log Y$ each. Use $e = b_0faq$, for which $e \ll Y^{1+\eta}$ and $(e, N) = 1$.

Sieve $w = (N - p)/e$ to $z = Y^\kappa$ on the new fixed interval for n/N, as in (11). Its length is a fixed positive multiple of N. Discard divisors $v \mid w$ in $[cY^{1-\eta}, CY]$ for fixed sufficiently small c and large C depending on the chosen packet. This costs $O(\eta/\kappa + o(1))$ times the same sieve scale; choose η last, also in terms of the interval length. The small part b_0fa and the marked prime q are unique after this discard. Remove nonsquarefree partners as before. Cancelling the two terms at b_0 leaves the divisors of $faqw$ in $(T/b_0, T]$. Any such divisor must contain q, since otherwise it uses a forbidden divisor of w, or is at most $fa \ll Y^\eta$. A divisor containing q cannot use w since $qz \gg Y$. By isolation, its a-part is exactly b. Its remaining f-part is therefore one of the designated divisors. Their signed sum at the ith location telescopes to c_i by the construction. This proves the lower bound. A nonzero entry is bounded above by the first absolute moment at its cutoff, by Theorem 1.

For the fixed shift the interval and the residue in BV become $p = n - d$ and $-d$ for the prime slot modulo e. The two-form sieve uses $u, evu - d$, with the same local factors at odd primes dividing dev. In the upper bound the two forms in (8) become $r, mqr - d$, and Henriot applies to $n, n - d$ with fixed discriminant. The complement uses $\log(X/\operatorname{rad} n)$ on $n \le X$; the at most $O_d(1)$ boundary values cost $X^{o(1)}$. All other estimates only use these one- and two-form bounds. This proves the shifted assertions. □

For instance, on a power-of-two total each fixed nonzero integer occurs with frequency of order $B(\log Y)$ and may be prescribed together with arbitrary values at finitely many other proportional cutoffs. The local constructions leave the multiplier $-\mu(n)$ on the large cutoffs. They do not estimate its mean.

Corollary 8. *For every fixed $0 < s < s_*$, uniformly as $Y = \min(R, X/R) \to \infty$,*

$$X^{-1} \sum_{n \le X} |M_R(n)|^s \asymp_s B(\log Y).$$

Proof. When $R \le X^{3/4}$ use Corollary 2, since then $\log Y \asymp \log R$. For $R > X^{3/4}$ use the construction of Theorem 4 with $Y = X/R < X^{1/4}$ on ordinary integers $n = ew \in (X/2, 3X/4)$. For each e an interval lower sieve gives $\gg X/(e \log z)$ choices of w, since its length is $\gg X/Y^{1+\eta}$ and the level z^4 is admissible with integer progression errors at most one. For a discarded z-rough v the single upper sieve on w/v costs $O(X/(ev \log z))$. The squarefree exceptions cost $O(X/z)$ by counting

multiples of prime squares. All the isolated-divisor estimates hold with mass $1/e$, already used in deriving (10). Thus the same complementary value of absolute size one occurs on $\gg XB(\log Y)$ integers. Theorem 1 supplies the upper bound. □

8 The same positive second-moment constant at complementary cuts

Fix an even squarefree integer D. As $N \to \infty$ through totals with $\operatorname{rad} N = D$, let $Y \to \infty$, $Y \le N^{1/4-\epsilon}$. With $M_R(n) = \sum_{d\mid n, d\le R} \mu(d)$, we have

$$\frac{1}{\pi(N)} \sum_{p<N} M_Y(N-p)^2 \longrightarrow C_D, \quad \frac{1}{\pi(N)} \sum_{p<N} M_{N/Y}(N-p)^2 \longrightarrow C_D. \tag{PM1}$$

The constant is strictly positive. It has a convergent series with nonnegative terms below. The same two statements hold for a fixed nonzero even shift $n = p + d > 0$, $p \le X$, with $D = \operatorname{rad}|d|$ and X in place of N. For integers the second-moment constant and the equality at complementary cuts are known. Dress–Iwaniec–Tenenbaum proved the limiting quadratic estimate and de la Bretèche–Dress–Tenenbaum treated the full two-parameter mean. The arguments here use the Selberg diagonalization, a prime-row change of local probabilities, and ordinary Bombieri–Vinogradov. Dress–Iwaniec–Tenenbaum · de la Bretèche–Dress–Tenenbaum · Selberg's diagonalization, chapter 21

8.1 The independent local model

For each prime $r \nmid D$ independently draw a nonnegative integer V_r with $\Pr(V_r \ge k) = 1/\varphi(r^k)$, $k \ge 1$. At $r \mid D$ set $V_r = 0$. Write $d \mid V$ when all its required valuations are present. In a divisor sum only primes through its cutoff matter. Put $g_r = \Pr(V_r \ge 1)$, $u_r = r g_r(1-g_r)$, and extend g multiplicatively on squarefree integers. For $r \nmid D$, $u_r = 1 - 1/(r-1)^2$. Put

$$B_D = \prod_r (1-1/r)(1+u_r/r), \qquad h_r = (1+u_r/r)^{-1}.$$

The positive number claimed in (PM1) is

$$C_D = B_D \sum_{j\ge 1} \log(1+1/j) \sum_{m,n\le j} \mu(m)g(m)\mu(n)g(n) \prod_{r\mid mn} h_r. \tag{PM2}$$

Each inner sum is the mean square of $\sum_{m\le j} \mu(m)g(m)1_{(m,K)=1}$ with independent events $r \nmid K$ of probabilities h_r. The first term gives $C_D \ge B_D \log 2 > 0$.

Expand the indicators $1_{r\mid V} = g_r + (1_{r\mid V} - g_r)$ and use orthogonality. For real $T, T' > 0$, the covariance without mean subtraction is exactly

$$\begin{aligned} \mathbb{E}[M_T(V)M_{T'}(V)] &= \sum_{k\le \min(T,T')} \frac{a(k)}{k} S_k(T/k)S_k(T'/k), \\ a(k) &= \mu^2(k) \prod_{r\mid k} u_r, \qquad S_k(t) = \sum_{m\le t,(m,k)=1} \mu(m)g(m). \end{aligned} \tag{PM3}$$

For fixed b, the mean of $a(k)1_{(k,b)=1}$ is $B_D \prod_{r\mid b} h_r$. Divide its Dirichlet series by zeta at 1, or apply convolution with 1 and absolute convergence of the residual coefficients with weight $1/n$. When $T = T'$, sum on $T/(j+1) < k \le T/j$ for fixed j to obtain each term of (PM2).

Here is a uniform tail estimate, also needed at a moving cutoff. It holds for any local probabilities with $g_r = 1/r + O(r^{-2})$ outside a finite excluded set F and zero on F, with a fixed absolute constant in the O-term in both applications. For any fixed A and $t \geq 1$,

$$|S_k(t)| \ll_A (1+\log t)^{-A} \prod_{r \mid k \prod_{p\in F} p} (1 - r^{-3/4})^{-1}. \tag{PM4}$$

Indeed the Dirichlet series factors as $1/\zeta(1+s)$ times local factors $(1 - g_r r^{-s})/(1 - r^{-1-s})$, with numerator one at the excluded primes and at $r \mid k$. Their absolute coefficient sum with weight $n^{1/4}$ has the displayed product bound. Convolve with $\sum_{m\leq t} \mu(m)/m \ll_A (1+\log t)^{-A}$, splitting the other variable at $\sqrt{t}$. The mean of $a(k) \prod_{r\mid k}(1 - r^{-3/4})^{-2}$ is bounded by an absolute constant uniformly in F, since its nonzero prime factors are $1 + O(r^{-3/4})$. Its sum to x is $O(x)$ by a positive convolution majorant. Thus the part of (PM3) with $k \leq T/J$, when $T = T'$, is $O_{D,A}((1+\log J)^{1-2A})$ by dyadic subdivision. This proves convergence and (PM2). It also proves that $\mathbb{E}[M_{cT}(V) M_{c'T}(V)]$ has a limit for every fixed $c, c' > 0$: below either fixed truncation partition by the integer parts of both arguments, then take a dyadic tail. The argument gives a uniform second-moment upper bound $O(\prod_{p\in F}(1 - p^{-3/4})^{-2})$ for all real cutoffs.

8.2 Transfer before and after squarefree complementation

Write $L = \log N$. At $R = Y$ the doubled divisor sum has progression moduli $[m, n] \leq Y^2$. Each occurs at most $\tau([m, n])^2$ times. On the reduced classes ordinary Bombieri–Vinogradov with a divisor weight gives the independent mean square with error $O_A(L^{-A})$, on normalizing the prime logarithm by N. For the divisor weight combine arbitrary logarithmic precision from BV, Cauchy's inequality and the bound $E_q \ll_\epsilon N/\varphi(q)$ from Brun–Titchmarsh at these shortened levels. Nonunit classes here, and below, contain only the finitely many primes dividing D; delete those primes before the expansion. The aggregate correction to reduced progression errors after deleting them is negligible.

Introduce $H(V) = \prod_r r^{(V_r - 1)_+}$; only finitely many primes repeat almost surely. For integer $n > 1$ put $h = n/\operatorname{rad} n$. Squarefree complementation gives the exact identity $M_{N/Y}(n)^2 = M_{<nY/(Nh)}(n)^2$. The analogous main term from prime progressions is

$$\int_0^1 \mathbb{E}\, M_{tY/H(V)}(V)^2\, dt + o(1). \tag{PM5}$$

To verify this before any limit in Y, put $B = L^K$ for a sufficiently large fixed K. Use $h \leq B$ and write $a = h \operatorname{rad} h$. The exact predicate for $n/\operatorname{rad} n = h$ is $a \mid n$, $(n/a, h) = 1$, and n/a squarefree. Expand the latter two tests over $b \mid h$, $j^2 \mid n/a$, with weights $\mu(b), \mu(j)$, keeping $j \leq B$. Errors are supported on numbers with a square divisor v^2 for $v > \sqrt{B}$. Their number is $O(N/\sqrt{B})$; the absolute error at one integer, including all false terms of the expansion and both divisor sums, is at most $O(\tau(n)^6)$. Cauchy with the elementary fixed divisor moments absorbs this after taking K sufficiently large. The same discarded terms have negligible model mass. There $\sum_{v>\sqrt{B}} \Pr(v^2 \mid V) \ll 1/\sqrt{B}$ by $\sum_{v>x} 1/\varphi(v^2) \ll 1/x$; the divisor function on the valuations up to N has any fixed moment bounded by a fixed power of L by its Euler product. A repeated prime above N has probability $O(1/N)$ and is independent of the small valuations. These arguments also cover the atom $n = 1$ separately.

For each remaining pair of sharp divisors $d, e \leq Y/h$, the modulus is $[d, e, a[b, j^2]] \leq Y^2 L^{O(K)}$. All five variables d, e, h, b, j divide it. The prime count has one variable endpoint imposed by $n/N > h \max(d, e)/Y$. Weighted BV with endpoint maximum therefore applies as at the short

cut. Its main term is the uniform t integral with independent valuations, because the probability of every allowed divisibility event at modulus q is $1/\varphi(q)$. Endpoint equality has zero measure in that integral. The prime logarithm may be removed in both comparisons: on $p \geq N/L^J$ it is $L + O_J(\log L)$; the shorter interval costs $o(N/L)$ in second moment by Cauchy and the elementary fourth divisor moment with sufficiently large fixed J. For fixed shifts the geometric variable is $(p+d)/X$; the endpoint translation is $O_d(1/X)$ and total over moduli contributes $X^{-1/2+o(1)}$ after normalization. Finitely many endpoint partners of size $O_d(X)$ cost $X^{o(1)}$.

To evaluate (PM5), condition on the full squarefull part $a = \prod_{V_r \geq 2} r^{V_r}$. Outside Da the remaining prime indicators are independent with probabilities

$$g'_r = \Pr(V_r = 1 \mid V_r \leq 1) = \frac{1/r}{1 - 1/(r(r-1))};$$

at primes dividing a they are forced to one. Expand over $d \mid \operatorname{rad} a$. By the two-cutoff limit following (PM4), the conditional mean square has a limit as its argument tends to infinity. That limit is the same for Y and $tY/H(a)$ for every fixed $t > 0$. Its uniform upper bound at all scales is

$$\ll_D 4^{\omega(a)} \prod_{r \mid a} (1 - r^{-3/4})^{-2},$$

by Cauchy and (PM4). This is integrable against the squarefull-part probabilities, bounded by $1/\varphi(a)$, since every exponent in a is at least two. Dominated convergence recombines the same limit C_D in both cases. This proves (PM1).

8.3 Concentration of the second moment

For a fixed prime radical D in the range of (PM1), almost all of either second moment comes from partners with $|M_R(n)| > (\log Y)^c$, for every $c < \log 2 - 1/2 = 0.193147\ldots$. Take fixed $s < 2$ sufficiently close to 2. The discarded mean square is at most

$$(\log Y)^{c(2-s)+2^{s/2}-1-s/2} = o(1).$$

The exponent follows from the small-moment upper bound. At the same time, on the support the conditional absolute moments of each order below 1.0575 remain bounded. Both statements concern the actual prime rows.

9 A Gaussian law for the size of the sharp Möbius divisor sum

Fix an even radical $D = \operatorname{rad} N$, write $Y \to \infty$, $Y \leq N^{1/4-\epsilon}$, $l = \log Y$, $U = \log l$, and take either $R = Y$ or $R = N/Y$. Sample $n = N - p$ with probability proportional to $M_R(n)^2$, where $M_R(n) = \sum_{d \mid n,\, d \leq R} \mu(d)$. Count distinct prime divisors through Y by $K(n) = \omega(n; \leq Y)$. the joint limit is

$$\left(\frac{K - 2U}{\sqrt{2U}}, \quad \frac{\log |M_R| - (\log 2 - 1/2)U}{(\log 2)\sqrt{U/2}} \right) \implies (Z, Z), \qquad Z \sim \mathcal{N}(0,1). \tag{GL1}$$

The singular limiting covariance has a quantitative explanation. With this sampling law,

$$2 \log |M_R| = (\log 2)K - U + O_{\mathbb{P}}(\log(2+U)). \tag{GL2}$$

The same conclusions hold for each fixed even nonzero shift $n = p + d > 0$, replacing N by X, and for ordinary integers at least throughout $Y \le X^{1/2-\epsilon}$. The normalization of the prime-row weights has a positive limit by (PM1). Granville–Koukoulopoulos–Maynard, Theorem 1.6(b), give the ordinary-integer square-weighted concentration of the number of prime factors around $2 \log\log R$ to any fixed relative tolerance. Their Theorem 1.3 gives the independent fourth-moment log power two. Their Theorem 10.2(b) already bounds the integer mean of $v^{\Omega(n;R)} M_R(n)^{2k}$ by $(\log R)^{v\binom{2k}{k}-2k} (\log\log R)^{O(1)}$ for $1 - 1/\binom{2k}{k} + \epsilon \cdot 1_{k=1} \le v \le 2 - \epsilon$. For $R \le X^{1/2}$, where $\log R = l$, at $k = 1$ this is (GL4) on integers for $1/2 < v < 2$, and at $k = 2$ the upper branch of (GL3) on integers for $5/6 \le v < 2$, in both cases up to the log-log factor and with $\Omega(n; R)$ in place of K; for larger R their bound is a power of $\log R$, not of $\log Y$. Chernoff applied to it, as in their proof of Theorem 1.6(b), gives the upper rate in (GL8) on integers for $2 < c < 4$ and $R \le X^{1/2}$. The new assertions here concern the joint Gaussian limit, the coupling to the actual size of the divisor sum, and transfer to individual prime rows. GKM, Theorems 1.3 and 1.6

Gorodetsky and Grimmelt proved the Gaussian count law on the reflected prime row with the full weight $2^{\omega(N-p)}$, settling Elliott's conjecture. That is the comparison measure suggested by (GL2). Here the sampling weight is the square of the sharp Möbius sum at one of two reciprocal cutoffs, and the mark stops at Y. The comparison alone does not identify the Gaussian under this change of measure; the marked quadratic asymptotic in (GL7) supplies that step. A proof in Section 10 extends (GL2), (GL8) and the fractional lower bound to every fixed exponent below half. Gorodetsky–Grimmelt, Theorem 1.1

9.1 The fourth moment with a factor-count weight

Let $W = X$ on integers and $W = \mathfrak{S}(N)N/\log N$ on an even prime row. The following bounds hold for $Y = \min(R, X/R) \to \infty$, throughout the full range of cutoffs. For fixed $0 < v \le 1$,

$$W^{-1} \sum_{\text{row}} v^K M_R^4 \ll_v \begin{cases} l^{-1}, & v < 1/2, \\ l^{-1}(1+U)^7, & v = 1/2, \\ l^{6v-4}, & v > 1/2. \end{cases} \tag{GL3}$$

For $v > 1/2$, in a fixed compact positive interval (also allowing $v > 1$),

$$W^{-1} \sum_{\text{row}} v^K M_R^2 \ll_v l^{2v-2}. \tag{GL4}$$

In particular the ordinary fourth-moment upper order holds on each individual prime row with its singular factor at both cutoffs.

Use the two-scale argument of Section 3 with its smooth bumps of height one and width $1/T$. Apply Hölder to the sum of large-factor translates with power four; the extra coefficient is $2^{3\omega(n;(z,Y^{5/4}])}$. Together with the choices of the translate this costs at most 16 per large prime. A fixed power of the first passing scale pays for these Euler factors as before. All discarded terms fit after increasing the stopping power to $T \asymp l^{1000}$ and the final Rankin constant. Indeed the eighth power of the crude divisor envelope costs $O(Wl^{255})$ before the weight; for any fixed larger local weight one may further increase the stopping power. Small primes $p \le z$ have weight v at each inclusion. When $v > 1$ the extra large prime weights up to Y cost another fixed power of the reciprocal trial exponent.

For a dyadic frequency block $|t| \asymp B$, integrate the fourth power of the small-list difference on the entire real line. Local Sobolev on the cells of size $1/T$ bounds its sampled fourth power by a constant times $\|F\|_4^4 + T^{-4}\|F'\|_4^4$ times the cell-density bound. This is uniform in the list-dependent

translation. Enlarge to all squarefree harmonic lists of primes through Y while each list's integrand is nonnegative. After Fourier inversion of this whole-line integral the four signed frequencies sum to zero (use conjugation for the last two). Independently summing each prime gives an Euler factor $1 + vb_p \prod_{j=1}^{4}(1 - e^{it_j x_p})$, where $x_p = \log p/l$ and the masses $b_p = C(p)/(p-1)$ may have the bounded Rankin inflation. Taking the absolute value of this product costs its real prime exponent plus $O_v(1)$, because $\sum b_p^2 \ll 1$ and the local functions are bounded. The change of a bounded total mass in the exponent also costs $O_v(1)$.

Put $V = \min(B, l)$. Primes with $x_p < 1/V$ contribute $O(1)$ for $B \le l$ by the product of four small differences, and for $B > l$ by the bounded mass of primes below e. On the remaining primes and for $|t| \le 8B$,

$$\sum b_p \cos(tx_p) = \log^+ \frac{V}{1+|t|} + O(1 + \log(1 + \log^+(B/l))).$$

This follows by partial summation on the physical log-prime bins, as in Section 2, and integration of cosine on $[1/V, 1]$. The trigonometric expansion of the real product has two constant terms. There are three complementary pairs of pair-sums, with positive signs. The eight odd-subset terms have negative signs and their real prime sums are bounded below by the error just given. Thus the product is bounded by

$$(1 + \log^+(B/l))^{O_v(1)} V^{2v} \prod_{j=2}^{4} \left(1 + \frac{V}{1 + |t_1 + t_j|}\right)^{2v}. \tag{GL5}$$

The three pair-sums are independent linear coordinates on $\sum t_i = 0$ with constant nonzero Jacobian. For $|t_1| \ge B$, at least one has absolute value $\ge 2B/3$, since their sum equals $2t_1$. Partition according to which, costing a factor three. On a box of side $O(B)$, the other two factors have one-dimensional mean $O_v(1)$ if $v < 1/2$, $O(1+\log V)$ at equality, and $O_v(1+V^{2v}/B)$ if $v > 1/2$. The lowest frequency block costs an absolute constant. Using three logarithms at equality is a harmless weakening.

The four Fourier envelopes and the Sobolev derivative therefore give the following bound on the fourth moment of the sampled block, apart from the arithmetic prefactor W/l:

$$T^{-4} B^3 (1 + B/T)^{-A} (1 + \log^+(B/l))^{C_v} V^{2v} \begin{cases} 1, & v < 1/2, \\ (1 + \log V)^3, & v = 1/2, \\ (1 + V^{2v-1})^2, & v > 1/2. \end{cases}$$

Here A may be arbitrarily large. The sampling uses full-line Fourier norms, so signs inside the Euler factor have been retained. Minkowski sums the fourth roots over B and T. For $T \le l$ the controlling powers per bump are T^{2v-1} below the transition and T^{6v-3} above it. For $T > l$ one gains geometric decay from T^{-1} with the other powers frozen at l. At $v = 1/2$ the $O(U)$ bumps and three logarithms give the stated exponent seven. The base smooth cutoff costs $O_v(W/l)$ by Fourier inversion and the multiplicative bound. This proves (GL3). For (GL4) use the squared Plancherel argument in Section 3 with effective mass $2v \log 2$ per octave, and no support-length factor. The total per bump is $(W/l) T^{-1} \min(T, l)^{2v}$ with the same harmless high-frequency logarithms. This gives (GL4) uniformly away from $v = 1/2$.

9.2 Coupling the size to the number of factors

Use (PM1) to normalize by a mass comparable to $N/\log N$. Write $A = lM_R^2/2^K$. Under square weighting,

$$\Pr(A \le e^{-t}) \ll_D e^{-t}, \qquad \Pr(A \ge e^t) \ll_D (1+U)^7 e^{-t}. \tag{GL6}$$

For the first inequality the numerator on this event is bounded by $e^{-t}2^K/l$; the row mean of 2^K is $O_D(l)$ by the multiplicative estimate. For the second use $l\,2^{-K}M_R^4e^{-t}$ and (GL3). This proves (GL2).

9.3 A tilted Selberg form and the count limit

First use the independent valuations of (PM2). Let $z \in [0.9, 1.1]$ be real and let $z \to 1$ as $Y \to \infty$. Weight by z^K. Its scalar normalization is

$$\prod_{p\le Y}(1+(z-1)g_p) = (1+o(1))l^{z-1}.$$

The remaining independent inclusion probabilities up to Y are $g_p(z) = zg_p/(1+(z-1)g_p)$. The ratio above tends to one because the bounded Euler remainder in its logarithm is $O_D(|z-1|)$. The same calculation holds for inclusion probabilities $1/p + O(p^{-2})$ with a fixed excluded set, as used after squarefull conditioning.

For any fixed $0 < c, c' \le 1$, the mixed form (PM3) with $g_p(z)$ is

$$l^{z-1}\{C(c,c') + o(1)\} \tag{GL7}$$

at cutoffs $cY, c'Y$, where $C(c,c')$ is precisely its unweighted limit. Probabilities above Y are irrelevant and can be extended by the same formula. To check (GL7), use $a_z(k) = \mu^2(k)\prod_{p|k} pg_p(z)(1-g_p(z))$. Its prime factors equal $z + O(1/p)$ outside the fixed excluded set. For every fixed integer b, the sum of $a_z(k)1_{(k,b)=1}$ up to x is $(C_z(b)+o(1))x(\log x)^{z-1}$ uniformly on the displayed z-interval; the leading factor is continuous in z. This is the classical Selberg–Delange theorem, factoring out $\zeta(s)^z$ with an absolutely convergent remainder on $\Re s > 3/4$. The same mean upper bound holds after multiplication by $\prod_{p|k}(1-p^{-3/4})^{-C}$ for any fixed C. The local remainder still converges absolutely on a half-plane strictly to the left of one. Thus every fixed cell in the argument of (PM3) gives its unweighted coefficient on division by l^{z-1}, in the limit $z \to 1$.

For the uniform tail, factoring the inner divisor series by $\zeta(1+s)^{-z}$ gives

$$|S_{z,k}(t)| \ll (1+\log t)^{-0.85}\prod_{p|k}(1-p^{-3/4})^{-C}$$

with a fixed constant depending on the excluded primes. The reciprocal-zeta-power partial sum at weight $1/n$ is $O((1+\log t)^{-0.9})$ uniformly in this z-interval. For example, apply the classical Selberg–Delange expansion to the coefficient sum and partial summation over the tail, using its convergent sum zero at 1. The factor left over has uniformly bounded absolute coefficient norm at shift 1/4 apart from the displayed excluded-prime product. This is the argument of (PM4). Dividing the $k \le Y/J$ tail by l^{z-1}, dyadic subdivision gives $O((1+\log J)^{-0.7}) + O(l^{-0.7})$. For $k \ge \sqrt{Y}$ use $\log k \asymp l$; below it use $\log(Y/k) \gg l$ and the mean upper bound summed with weight $1/k$. This proves (GL7). Selberg–Delange inputs

Put $z = \exp(v/\sqrt{2U})$ for each fixed real v. Combining (GL7) with the scalar normalization and centering gives

$$l^{2z-2}e^{-v\sqrt{2U}} \;\longrightarrow\; e^{v^2/2}.$$

This proves a standard Gaussian limit for $(K-2U)/\sqrt{2U}$ under the independent law weighted by M_Y^2, by convergence of moment generating functions. For the model of the complementary cutoff (PM5), condition on the full squarefull part a and on fixed $0 < t < 1$. Expand the fixed forced primes exactly as in (PM5). The mixed short cutoffs are fixed positive multiples of Y. Their

kernels have (GL7) with the squarefree-conditional local probabilities. The finitely many forced primes change the scalar normalization by $1+o(1)$. Hence each fixed conditional square form with positive limiting mass has the same Gaussian count law. Those with zero limiting mass disappear. Integrate the conditional *characteristic functions*, bounded by the unweighted conditional mass; its integrable majorant is given after (PM5). The resulting total has limit $C_D e^{-t_0^2/2}$ at frequency t_0.

9.4 Transferring a count mark to genuine primes

Take $m = \lceil C(1+U) \rceil$, $y = Y^{1/(C'(1+U))}$, with C sufficiently large and then C' sufficiently large depending on C and the quarter-level margin. First mark only $K_y = \omega(n; \le y)$. For $|\zeta| = 1$ near 1 expand ζ^{K_y} over squarefree products of at most m of these primes. The truncation error vanishes unless $K_y > m$, and in absolute value is at most $1 + 2^K$ there. Against M_R^2, its normalized integral is $o(1)$ both on the arithmetic and the model rows. Indeed $|M_R| \le 2^K$ at the short cut, also after squarefree complementation, and $1_{K>m} \le 2^{K-m}$. The multiplicative estimate gives $\mathbb{E} 16^K \ll_D l^{15}$; take C sufficiently large. The model has the same elementary bound. The transfer calculation (PM5) now only enlarges the modulus by y^m. Each coefficient still has absolute value at most one and at most one extra divisor of the modulus is specified. Both the squarefull truncation and divisor-weighted BV apply, since $Y^2 y^m L^{O(1)} < N^{1/2-c_\epsilon}$. For squarefull truncation the marked polynomial is bounded by $2^K \le \tau(n)$; increase the fixed logarithmic cutoff if needed. Thus the marked arithmetic and model characteristic integrals differ by o(1), including the moving endpoint.

In either measure weighted by M_R^2, the omitted count $K - K_y = O_{\mathbb{P}}(\log(2+U))$. Indeed the unnormalized second moment with extra factor 2^{K-K_y} is $O_D((1+U)^C)$ in units $N/\log N$. In the squared-bound proof, the extra prime harmonic mass on $(y, Y]$ is $O(1 + \log(1+U))$. Its upper effect on the Euler exponential, both large and small lists, is a factor $(1+U)^{O(1)}$. All local weights remain bounded and the previous polynomial remainders still fit. For the independent model the identical bound follows also by taking the prime-row limit at each **fixed** Y and its associated y: the finite prime mark has an exact divisor expansion and the squarefull truncation in (PM5) applies before letting Y grow. Thus replacing K_y by K at frequencies $t_0/\sqrt{2U}$ costs o(1); the difference of the phases is bounded by two. The same proof on integers uses ordinary counting errors at most one for each modulus, with total $O(Y^2 y^m X^{o(1)}) = o(X)$ below the stated half level. Together with (GL2) this proves (GL1).

9.5 A lower bound on fractional moments above the support scale

For each fixed $1 < s < 2$ in the range of (GL1), put

$$F(s) = (4-s)\{2^{-(2-s)/(4-s)} - 1\} + (2-s)/2.$$

On either genuine prime cutoff the unweighted mean of $|M_R|^s$ is at least $l^{F(s)-o(1)}$. It strictly exceeds the support scale in its power of l for $s > 1.38283672$. To see this put $z = 2^{-(2-s)/(4-s)}$ and use Hölder:

$$\mathbb{E}|M_R|^s \ge \frac{(\mathbb{E} z^K M_R^2)^{(4-s)/2}}{(\mathbb{E} 2^{-K} M_R^4)^{(2-s)/2}}.$$

The denominator uses (GL3). For the numerator one needs $\mathbb{E} z^K M_R^2 \ge l^{2z-2-o(1)}$ for fixed $0.79 < z < 1$. Use the preceding truncated count expansion for K_y, now a lower Bonferroni polynomial of odd degree m. Its model error is $O(l^{-10})$ by the same elementary bound after increasing C. On the complementary row restrict to squarefree partners with $n/N \in [1/2, 3/4]$. In the model this costs a fixed squarefree density, with the modified inclusion probabilities of (PM5). For both cuts the

scalar tilt on primes through y contributes $l^{z-1+o(1)}$. In the remaining quadratic form keep only the terminal terms $T/2 < k \le T$, where $T = Y$ or tY. The inner sums equal one. The tilted or untilted local coefficient $pg'_p(1-g'_p)$ is at least a fixed positive weight $z + O_z(1/p)$ on each allowed prime (take the smaller of the two actual coefficients). The positive Selberg–Delange mean on this dyadic interval contributes $\gg_D l^{z-1}$. This proves the claimed lower bound for the shortened count. To put back $K - K_y$, its part above $AU/\log(2+U)$ has negligible second moment for a sufficiently large fixed A. If $l \le \sqrt{L}$, instead take $y = Y$, with $Y^{m+2} < N^{1/2-c}$, and there is no tail. If $l > \sqrt{L}$, the integer count with at least j distinct primes in $(y, Y]$ is at most $N(\sum_{y<p\le Y} 1/p)^j/j!$. Take $j = \lceil AU/\log(2+U)\rceil$, use Cauchy and the elementary fourth moment of 2^K; the cost of removing primality in this estimate is at most $L \le l^2$. The resulting error is $O(l^{-10})$ on increasing A. Below this threshold the extra factor $z^{K-K_y} \ge l^{-o(1)}$. The numerical threshold is the root of $F(s) = -\delta$ in (1,2). The final Section 11 improves this interpolation to the matching log exponent for every $s_* \le s \le 2$.

For completeness (GL3) and (GL4) also improve interpolation **above** order two on arbitrary individual prime rows. Write $s = 2 + a$, $0 < a < 2$, $v = \max(1/2, 3^{-(2-a)/2})$, $w = v^{-a/(2-a)}$. Hölder on $(v^K M_R^4)^{a/2}(w^K M_R^2)^{1-a/2}$ gives a log exponent

$$\tfrac{a}{2}\max(-1, 6v-4) + (2-a)(w-1).$$

At $v = 1/2$ there is also a power of $1+U$. In particular the upper exponent of the third absolute moment is $2\sqrt{3} - 3 = 0.4641016\ldots$ and the fourth-moment exponent is two. The exponents from both sides have derivative $\log 2 - 1/2$ at order two, as does the lower fractional bound. On the left the latter and the upper bound have the same quadratic Taylor term $(\log 2)^2(s-2)^2/4$. The finite lower threshold above is unique because $F(s) = (4-s)(2^{2/(4-s)-1} - 1/2) - 1$ is increasing for $1 < s < 2$.

9.6 Full large-deviation rate for the number of factors

Under the same square weighting, K/U satisfies a large-deviation principle of speed U. Its good rate function is

$$I(c) = \begin{cases} +\infty, & c < 0, \\ 1 - c\log 2, & 0 \le c \le 1, \\ c\log(c/2) - c + 2, & c \ge 1. \end{cases} \tag{GL8}$$

The linear part needs a separate lower construction because the scalar moment exponent has a corner. Here all logarithmic probabilities are in units of U.

For $0 < a \le 1$ restrict the partner to have no prime factor at most $w = \exp(l^{1-a})$. Let $Z_a(z)$ be the original (unweighted row) mean of $1_{P^-(n)>w} z^K M_R^2$. For each fixed $z > 1/2$,

$$Z_a(z) = l^{-1+a(2z-1)+o(1)}. \tag{GL9}$$

For its upper bound, omit the primes through w in the small-list Plancherel procedure. They are below every trial cutoff for sufficiently large Y. As in the primorial calculation the number of effective frequency octaves is capped at $l/\log w = l^a$, since the final bin then has prime mass $O(1)$. The outer factor is still W/l. The squared norm per bump has multiplier $T^{-1}\min(T, l^a)^{2z}$ before the harmless high-frequency logarithms. This proves the upper bound. With $a = 1$ we can also retain all primes. For $z \le 1/2$ the all-prime version of this sum has upper order $l^{-1+o(1)}$ by the same geometric sum, logarithmic at equality. In this exponent range the matching lower bound comes from the partners with no prime through Y, by the positive lower linear sieve at the

shortened quarter level. On central squarefree partners both divisor cuts have magnitude one. The nonsquarefree exceptions for these rough partners are negligible by the square-divisor estimate.

For the lower bound in (GL9), transfer the local mark up to y as before, setting its prime factors equal to zero through w and equal to z on $(w, y]$. For fixed z the degree-m polynomial remainder vanishes when $K_y \le m$ and is bounded by $(2+2z)^K$ in general. The same exponential cutoff, with C and its successor C' enlarged depending on z, makes the normalized error $O(l^{-10-10z})$. Either discard $K - K_y$ beyond $O_z(U/\log(2+U))$ as in the preceding section when $z < 1$, or retain the inequality $z^K \ge z^{K_y}$ when $z \ge 1$. In the model the scalar normalization is

$$l^{-1}(l/\log w)^{z+o(1)}.$$

Indeed excluding the initial primes costs $\asymp 1/\log w$, then the tilt up to y costs $(l/\log w)^{z-1+o(1)}$. If necessary restrict to all squarefree valuations as at the complementary cut, which costs a fixed factor and changes the prime probabilities only by $O(p^{-2})$.

It remains to lower-bound the tilted quadratic form by $(l/\log w)^{z-1+o(1)}$. For a cutoff $T \asymp Y$, keep $T/w < k \le T$ in (PM3): the inner sum is exactly one because all primes through w are excluded. Write $k = mr$ with a prime $r \in (T/(wm), T/m]$ and a squarefree $m \le Y^{1/4}$ supported on $(w, Y^\eta]$, with $\eta > 0$ fixed small depending on z. The prime r exceeds $Y^{1/2}$, so the parametrization is unique. Its harmonic mass is $\gg_z \log w/l$, by the PNT on fixed-ratio intervals and summation. The harmonic mass of m with its local variance weights is $\gg_z (l/\log w)^{z-o(1)}$. On primes $\le y$ the factor is $z/p + O_z(p^{-2})$; the primes above y up to Y^η can simply be omitted, losing at most a factor $(1+U)^{O_z(1)}$. To impose the cut on m, its expected logarithm in the normalized squarefree Euler product is $O_z(\eta l)$. Markov therefore retains at least half on decreasing η. All primes dividing the fixed radical can be removed at bounded cost. These estimates prove (GL9), also on the fixed-shift and integer rows.

For the full row the logarithmic generating function of the square-weighted count, divided by U, tends to $\max(-1, 2e^t - 2)$ at every real parameter t. Chernoff's bound gives the upper rate in (GL8), including exponential tightness by any positive t. For the lower rate, under the restricted weighting $1_{P^- > w} z^K M_R^2 / Z_a(z)$, (GL9) at $ze^{\pm h}$ implies $K/U \to 2az$ in probability, for fixed $z > 1/2$. On returning to square bias, every fixed open neighborhood of $c = 2az$ has lower log probability at least $-1 + a(2z-1) - c\log z$, by first restricting to arbitrarily small neighborhoods of c. For $c > 1$ choose $a = 1, z = c/2$. For $0 < c \le 1$ let $z \downarrow 1/2$ and $a = c/(2z)$. Finally at zero take the genuinely Y-rough partners, or take positive c down to zero in an open neighborhood. This yields the whole lower rate, including the straight segment.

10 A second-moment lower bound up to the half level

This extends the order estimates and the large-deviation principle from the quarter level to every fixed level below one half on genuine reflected primes. Fix D even and $\operatorname{rad} N \mid D$, and let $Y \to \infty$, $Y \le N^{1/2-\epsilon}$. Put $l = \log Y$, $U = \log l$, and take $R = Y$ or $R = N/Y$. With $K = \omega(N-p; \le Y)$, for each fixed $z > 1/2$ and $0 < a \le 1$, let $w = \exp(l^{1-a})$. Then

$$\frac{1}{\pi(N)} \sum_{p<N,\, P^-(N-p)>w} z^K M_R(N-p)^2 \asymp_{D,z,\epsilon,a} l^{-1}(l/\log w)^{2z-1}. \tag{MV1}$$

A bounded initial w can be omitted. The lower bound holds in the central interval on squarefree partners. The same argument applies on every fixed even difference. In particular at all direct exponents strictly below half, and at their reciprocals, the second moment stays between two fixed

positive constants. The common limiting constant and the Gaussian shape in (PM1) and (GL1) are proved through the quarter level. The full large-deviation rate (GL8), the coupling (GL2), and the fractional-moment lower exponent $F(s)$ now hold throughout the larger range. Section 11 uses (MV1) there and identifies the full upper branch in the log exponent.

The construction uses an interior marked prime and a core of size at most Y^{η}. Instead of selecting one isolated divisor, integrate the **square of the complete signed short divisor sum** of that core. This is the prime-row analogue of the harmonic lower-bound strategy in GKM §9. On the moving long cut the linear sieve is applied on separate intervals in the variable endpoint; one modulus has at most a divisor-bounded number of intervals. GKM §9

10.1 The marked window

Use the upper bound in (GL9), whose proof is uniform at all complementary cutoffs. Write $B = \max(20, w)$; if needed increase its fixed lower bound in terms of D. Choose a prime $b \in (B, B^2]$ outside D. Let $m \le Y^{\eta}$ be squarefree with primes all larger than B^2 and coprime to D. Here $\eta > 0$ is chosen sufficiently small after a fixed $\kappa > 0$ as in the proof of Theorem 4. Given a prime q, impose $n = N - p = bmqv$ and $P^-(v) > Y^{\kappa}$. Use q in $[cY/(bm), Y]$, with a fixed $c < 1/3$. Discard v admitting a divisor in $[cY/(B^2Y^{\eta}), Y]$. The dimension-one lower sieve and the dimension-two upper sieve for this discard are those in Section 5. On each retained partner there is exactly one possible b and one q: b is the smallest prime and any second candidate for q would divide v in the excluded window. Every prime factor of m lies below Y^{κ}.

On the direct cut use the fixed window $T = Y$; on the complementary cut use $T = tY$ with $t = n/N \in [1/2, 3/4]$ and restrict to squarefree n. On the retained rows,

$$M_T(n) = \mu(q) J_{m,b}(\log(T/q)), \qquad J_{m,b}(x) = \sum_{d \mid m,\ e^x/b < d \le e^x} \mu(d). \tag{MV2}$$

Use open endpoints at the complementary cut. Indeed cancel at b. A divisor in $(T/b, T]$ avoiding q would need the discarded divisor of v; m itself is too small. One containing q cannot use any prime of v since $qY^{\kappa} > Y$. On squarefree partners complementation changes only a unit sign. Also discard cofactors v with more than a fixed large number of prime factors in $(Y^{\kappa}, Y]$. On each allowed interval, counting one such prime costs at most the upper-sieve scale times $O(\sum_{Y^{\kappa} < r \le Y} 1/r) = O_{\kappa}(1)$. Use the two linear forms after extracting r. Markov makes this discard smaller than one fourth of the lower main term, uniformly in the fixed modulus after summing its BV errors. This ensures $z^K \gg_z z^{\omega(m)}$, with a constant depending also on ϵ, D.

For details at the moving cut, at fixed b, m, q split $t \in [1/2, 3/4]$ at the changes of (MV2). The number of resulting intervals is $O(\tau(m))$; the squared weight on each is at most $\tau(m)^2$. Omit intervals of length below $Y^{-3\eta}$. The lower sieve has main term proportional to the length and modulus $bmqd \le Y^{1+2\eta+4\kappa} < N^{1/2-c_{\epsilon}}$, with a BV discrepancy uniform in the two endpoints. All multiplicities and the additional factor $z^{\omega(m)}$ are at most a fixed power of the divisor function of the combined modulus. For the two upper-sieve discards the remaining integer ranges have length $\gg NY^{-3\eta}/(bmqr)$, or the same with a forbidden divisor in place of r; this is $\gg N^{c_{\epsilon}}$ on shrinking η. Both upper sieves are proportional to their respective lengths, with constants uniform in the retained cells. The forbidden-divisor harmonic cost is $O((\eta + \log B/l)/\kappa)$. Squarefree exceptions after restoring distinct partners have negligible weighted mass by Cauchy: a repeated prime then exceeds Y^{κ}, apart from the already excluded possibility $q \mid v$, and the count is $O_D(N/(LY^{\kappa}) + N^{3/4})$, $L = \log N$. The crude fixed multiplicative moments of $(2 + 2z)^K$ grow only like fixed powers of l. The direct cut only needs the original fixed central interval.

10.2 Plancherel on the small core

For every $t \in [0.56, 0.69]$, or for the direct cut $t = 1$, summing over the marked primes gives

$$\sum_q \frac{|J_{m,b}(\log(tY/q))|^2}{q} 1_{\text{retained interval}} \gg \frac{1}{l} \int_{\mathbb{R}} |J_{m,b}(x)|^2 dx. \tag{MV3}$$

Its support is $[0, \log(bm)]$, up to endpoints. The different knots are logarithms of distinct integers d, bd at most bm, so spacings are $\gg 1/(bm) \gg Y^{-2\eta}$. Remove physical log margins $10Y^{-3\eta}$ at both ends of each gap. This retains a fixed fraction of every gap. On its remaining q interval the t-cell reaches farther than $Y^{-3\eta}$ in both directions, or the direct cut is fixed. The q interval at height $Q \gg Y^{1-2\eta}$ has length at least $cQY^{-2\eta}$, hence exceeds $Q^{9/10}$ after choosing η small. On each such interval the prime number theorem in short intervals gives reciprocal mass bounded below by a constant times its logarithmic length divided by l. For larger ratios, subdivide into multiplicative intervals of bounded ratio, combining any final short piece or using a uniform geometric partition. Only an exponent strictly below one is needed in this use of Hoheisel's theorem. Summing the nonnegative constant heights proves (MV3). Huxley's short-interval estimate

Consequently the normalized weighted prime count is bounded below by a constant times

$$\frac{1}{l^2} \sum_{B<b\le B^2}^{\text{prime}} \frac{1}{b} \sum_m \frac{z^{\omega(m)}}{m} \int |J_{m,b}(x)|^2 dx. \tag{MV4}$$

One factor of $1/l$ is from the cofactor lower sieve, the other from q. The replacement of $1/\varphi(bmq)$ by $1/(bmq)$ is a lower bound. In the moving case integrate (MV3) over its displayed fixed t interval after applying the progression estimates.

Plancherel evaluates the last integral as

$$\frac{1}{2\pi} \int_{\mathbb{R}} \left| \frac{1 - e^{i\xi \log b}}{\xi} \right|^2 \prod_{r|m} |1 - r^{i\xi}|^2 \, d\xi.$$

Restrict to $1/\log B \le \xi \le 3/(2\log B)$. Then $1 \le \xi \log b \le 3$ and the first factor is $\gg (\log B)^2$. Use m with primes only through $Y^{\eta'}$, $\eta' > 0$ chosen sufficiently small in terms of η, z. Its unrestricted harmonic Euler product is $\asymp_{D,z,\eta'} (l/\log B)^{2z}$, uniformly in this range of ξ. Indeed $\sum_{B^2<r\le Y^{\eta'}} \cos(\xi \log r)/r = O(1)$ by the PNT integral and integration by parts from $2\xi \log B \ge 2$; the sum of prime reciprocals is $\log(l/\log B) + O_{\eta'}(1)$. The expected log m under the normalized nonnegative product is $O_z(\eta' l)$, so the bound $m \le Y^{\eta}$ retains at least half. Integrating the frequency interval and summing $1/b \gg 1$ over the allowed primes shows (MV4) is $\gg l^{-2}(\log B)(l/\log B)^{2z}$, which is (MV1).

For $z \le 1/2$ in the larger range it is unnecessary to lower-sieve Y-rough partners at Y itself. In the unmarked count exponent use (MV1) at a small fixed a and a tilt $z_0 > 1/2$ to get small positive counts; let a decrease after the large-N limit. This gives the lower rate at c=0 in each open neighborhood, and by the same tilting also the exponent -1 of the generating function. Together with (GL4) its large-deviation consequences follow exactly as in (GL8). At $a = 1$ the constants in (MV1) stay bounded above and below as z ranges on a compact neighborhood of one. Thus the count has centered fluctuations $O_{\mathbb{P}}(\sqrt{U})$ throughout the half-level range by Chernoff at $z = \exp(\pm A/\sqrt{U})$; the same is true for the log amplitude after (GL2). The normal limit in the interval between the quarter and half levels is open here.

11 The power of log Y between the first and second moments

On each fixed-radical reflected prime row, at either $R = Y, N/Y$, for $Y \to \infty$, $Y \le N^{1/2-\epsilon}$, and $s_* \le s \le 2$,

$$\frac{1}{\pi(N)} \sum_{p<N} |M_R(N-p)|^s = (\log Y)^{\,2^{s/2}-1-s/2+o(1)}, \qquad s_* = 2\log_2(1/\log 2). \tag{FS1}$$

The same conclusion holds for each fixed even nonzero shift in the displayed range, and on ordinary integers throughout $Y = \min(R, X/R) \to \infty$, including the square root by (FS5). The endpoint's log-log factor is unresolved. For $0 < s < s_*$ the sharper two-sided estimate already includes the log-log power $-3/2$. Thus s_* is the sharp boundary in the power of log Y between the support order and larger fractional moments. The argument below strengthens the interpolation lower bound in (GL7)–(GL9). It uses the second-moment lower estimate of (MV1) and a fourth moment restricted to the factor counts on several logarithmic intervals.

11.1 Counts under the weighted square

Write $l = \log Y$, $U = \log l$, set $\gamma = 2^{s/2} \in [1/\log 2, 2]$, and put $z = \gamma/2 > 1/2$. All arithmetic expectations below are normalized by $\pi(N)$, or by X on integers. The fixed singular factor is harmless. Let $K = \omega(n; \le Y)$. Divide the range of $\log p$ into J blocks

$$l^{i/J} < \log p \le l^{(i+1)/J}, \qquad 0 \le i < J,$$

and call their counts K_i. Factors with $\log p \le 1$ have bounded count. Fix J and a count tolerance $d > 0$. The set Γ given by $|K_i - \gamma U/J| \le dU/J$ for every i carries a fraction tending to one of the weight $z^K M_R^2$.

Indeed assign weights $z_p = ze^h$ in any one block and z elsewhere, with fixed small $|h|$ and $\min z_p > 1/2$. The squared Plancherel calculation of (GL4) gives

$$\mathbb{E}\left[M_R^2 \prod_{p\mid n, p\le Y} z_p \right] \ll_{J,z,h} l^{-2} \exp\left(2 \sum_{p\le Y} z_p/p \right). \tag{FS2}$$

For a bump width $1/T$ and a frequency block B, its squared Fourier amplitude integral is $O(B/T^2)(1 + B/T)^{-A}$. At $B \le l$ the effective Euler exponent is at most $2\sum_{p\le Y} z_p/p - 2\min(z_p)\log(l/B) + O_J(1)$. Split the oscillatory prime sums at the block endpoints. For $\log p < l/B$ the difference is small and contributes a bounded sum; above this threshold partial summation bounds each cosine integral by a constant. At $B > l$ the earlier high-frequency logarithms are summable. The outer arithmetic factor is $1/l$; summing squared roots over frequency scales and bump scales proves (FS2) since $2\min z_p > 1$. The smooth part is $O(1/l)$. The actual small-prime list in the sampling argument goes up to the trial cutoff; positivity of the integrated squares allows enlarging it to all primes through Y with the specified weights. The remaining large-prime Euler factors have only a fixed power of the reciprocal trial exponent. Now (MV1) without an initial exclusion bounds the untilted denominator below by a constant times l^{2z-2}. Use (FS2) at h and $-h$, and $\sum_{p\in I_i} 1/p = U/J + O(1)$, in Chernoff's inequality. This proves the assertion about Γ. In particular, with

$$D_2 = -2 + \gamma(1 + \log(2/\gamma)),$$

one has $\mathbb{E}[1_\Gamma M_R^2] \gg l^{D_2 - O_z(d)}$.

Keep only the original dyadic step bumps with $l^{1-\alpha} \le T \le l^{1+\alpha}$, for a small fixed $\alpha > 0$, and call the resulting divisor statistic F. At the complementary cut use the moving step and restore its actual complementary unit sign. By the same squared-block bound at constant z,

$$\mathbb{E} z^K |M_R - F|^2 \ll_z l^{2z-2-c_z\alpha}$$

for a fixed $c_z > 0$, also absorbing the negligible stopping window. Indeed the summed squared scale grows as T^{2z-1} on the lower side and decays as l^{2z}/T on the upper. Unweighting on Γ costs $l^{-\gamma \log z + O_z(d)}$. Hence these discarded terms have negligible square mass there if d is sufficiently small compared with α.

11.2 The fourth power on the retained counts

Put $D_4 = -3 + \gamma(1 + \log(4/\gamma))$. we obtain

$$\mathbb{E}[1_\Gamma |F|^4] \ll_{\alpha,J,d} l^{D_4 + O(\alpha + 1/J + d) + o(1)}. \tag{FS3}$$

Use the four-frequency calculation of (GL5) after applying the full-line cell Sobolev bound. The only blocks of prime sizes on which we impose counts during this calculation are $1 \le i \le J-2$. For fixed J these lie entirely below the smallest passing arithmetic cutoff $Y^{c/U}$ for large Y. Thus their counts are exactly the counts in the small list: the marked primes and all large sublists are outside them. Drop the other count restrictions when the term is still an integral of a nonnegative fourth power. The model sum may be extended to **all** squarefree lists in the first $J-1$ blocks that obey the retained counts, and freely elsewhere. Euler inflations from earlier trial failures change a bounded total mass. Abandoning the two end blocks costs $O(U/J)$ in all exponents below.

All four frequencies have size comparable to the same dyadic B, except in the bounded lowest block. Write them as t_i with sum zero, and use $b_j = t_1 + t_j$, $j = 2, 3, 4$, as independent coordinates. Keep $B \le l^{1+2\alpha}$. Above this point, $T \le l^{1+\alpha}$ and arbitrarily rapid Fourier decay bounds the remainder by $O(l^{-10})$ even with the *unrestricted* fourth Euler bound, by taking the smooth decay exponent sufficiently large depending on α. This discarding can also be done on the joint sampling norm before the coefficient calculation. Put each $1 + |b_j|$ in a dyadic interval of size Q_j. There are $O((1+U)^3)$ boxes. Their volumes are $O(\prod_j Q_j)$. Parameterize prime sizes by $x = \log\log p/U \in [0, 1]$. At this x let $r(x)$ count the three inequalities $Q_j l^{x-1} < 1$. The volume divided by l^3 is at most

$$O(1)\, l^{-\int_0^1 r(x)dx + O(\alpha)},$$

since each $Q_j \le O(l^{1+2\alpha})$. The determinant of the linear coordinate change is constant.

For each frequency quadruple, except for a bounded number of prime-size blocks containing or adjacent to a transition $Bl^{x-1} = 1$ or $Q_j l^{x-1} = 1$, the sum of the real local product $P_p = \prod_{m=1}^4 (1 - e^{it_m \log p/l})$, with harmonic prime weights, equals $(2 + 2r)U/J + o(U)$ if Bl^{x-1} is large. There $r \le 2$, since the three pair-sums add to $2t_1$. If Bl^{x-1} is small, the sum is $o(U)$ from the four small differences and is also bounded in absolute value by $2U/J + o(U)$. To verify the first estimate expand the trigonometric product. In a high block the four single and complementary triple terms oscillate. A small pair-sum gives two constant terms up to bounded total error; a large one has bounded integral. Partial summation by PNT is uniform above the first block because the physical log primes exceed $l^{1/J}$ while the frequency in physical log coordinates is at most $O(l^{2\alpha})$. The adjacency convention absorbs fixed dyadic factors even when a transition lies arbitrarily close to a block boundary. Empty transitions and the lowest frequency cost no additional blocks.

For clarity, for exactly n primes of one block take the coefficient of v^n in $\prod_p(1+vb'_pP_p)$, where b'_p are the harmonic masses with the bounded local corrections of the sampling proof. At radius $u>0$ its absolute value is at most

$$u^{-n}\exp\left(u\left|\sum_p b'_pP_p\right|+O_u(1)\right). \tag{FS4}$$

Indeed $\log|1+w|\le \Re w+|w|^2/2$, and $\sum_p(b'_p)^2\ll 1$. Apply this on the circle before the complex Cauchy coefficient bound. Replacing each local product separately by its absolute value would lose the rank calculation. For the nontransition constrained blocks choose $u=\gamma/A$, where $A=2+2r$ on the high blocks and $A=2$ on the low ones. Summing over allowed counts costs a factor $O(1+U)$ per block. Their contribution to the exponent per unit x, **after subtracting the volume cost**, is bounded by

$$\gamma(1+\log(A/\gamma))-r+O(d).$$

On high blocks the maximum over $r=0,1,2$ occurs at one: $\gamma\log 2\ge 1$ and $\gamma\log(3/2)<1$. On low blocks the same bound follows from $A=2$, $r\ge 0$. On transition blocks take radius one if extracting a coefficient and bound the Euler exponent by $O(U/J)$ per block using $|P_p|\le 16$. The same absolute bound treats the unconstrained end blocks. Since there are a bounded number, their total cost in the l-exponent is $O(1/J)$.

Finally the volume factor l^3, the Fourier envelopes $T^{-4}\le l^{-4+4\alpha}$, and the outer arithmetic factor $1/l$ cost $l^{-2+O(\alpha)}$. The sampling derivative factor $(1+B/T)^4$ costs at most $l^{O(\alpha)}$. The finitely many physical count blocks, dyadic frequency boxes, frequency blocks and retained bumps contribute $l^{o(1)}$ (Minkowski takes fourth roots before summation). Summing the first-passing scales works as in (GL3). This proves (FS3).

11.3 Returning to the sharp sum

Write $c=(D_4-D_2)/2=(\gamma\log 2-1)/2$. For any fixed $e>0$, choose $\alpha,1/J$ sufficiently small in terms of e, then take d still smaller in terms of both α and e. On $\Gamma\cap\{|M_R|>l^{c+e}\}$, either the discarded difference carries at least half the amplitude or $|M_R|\le 2|F|$. In the second case bound its square by $16l^{-2c-2e}|F|^4$. By (FS3) and the square estimate for the discarded difference, this removes a vanishing fraction of the lower square mass on Γ. Therefore

$$\mathbb{E}|M_R|^s\gg l^{D_2-O(d)-(2-s)(c+e)}.$$

Let e decrease. At $\gamma=2^{s/2}$, $D_2-(2-s)c=2^{s/2}-1-s/2$. The upper bound is already given by Theorem 1. This proves (FS1). At $s=s_*$ both powers equal $-\delta$, and (FS1) alone does not choose between the known log-log powers $-3/2$ and $-1/2$.

11.4 Ordinary integers at the square root

For ordinary integers (FS1) in fact holds whenever $Y=\min(R,X/R)\to\infty$, including $R=\sqrt{X}$. Only the lower bound for the tilted square in the range $X^{1/3}\le Y\le\sqrt{X}$ requires an additional argument. For smaller Y use (MV2)–(MV4) on integers: in its progression lower bound the count of multiples has error at most one at each modulus, and the cofactor upper sieves have one rough integer form. For $Y\le X^{1/3}$ the products in both discard estimates stay below $X^{3/4}$ on choosing the small exponents there sufficiently small. In particular the lower sieve and the discard both have

their respective length divided by $\log(Y^\kappa)$ as scale. We now show in the missing range that, for fixed $z > 1/2$,

$$X^{-1} \sum_{n \le X} z^K M_R(n)^2 \gg_z l^{2z-2}. \tag{FS5}$$

Fix a prime b, say 11. Choose squarefree $m \le Y^\eta$ with every prime greater than b. Sample integers $n = bmqr \in [X/2, 3X/4]$ with q and r distinct primes larger than $Y^{1-2\eta}$. The representation has multiplicity at most two: the small part is bm, and q and r can exchange places. Use $T = Y$ at the direct cut, $T = tY$ at the reciprocal cut, where $t = n/X$; in the latter case use squarefree complementation and the open step. Cancellation at b gives in absolute value

$$|M_R(n)| = |J(\log(T/q)) + J(\log(T/r))|, \qquad J = J_{m,b}.$$

Indeed the divisors of m alone occur below T/b, and $qr > T$. Put $u = \log(T/q)$, $H = \log(T^2/(tX))$. The second argument is $\log(bm) - u + H$. Reflection of the divisor pairs of m gives $J(\log(bm) - v) = \mu(m)J(v)$ off the knots. Thus the continuous integral of the square at fixed t is $\int |J(u) + \mu(m)J(u - H)|^2 \, du$. A nonzero summand in this integral forces q or r into $[T/(bm), T]$. As $Y^2 \le X$, both formal prime variables are then strictly above $2Y^{1-2\eta}$ and below $X/2$ for large Y and central t.

Here are details transferring the integral to primes. For the direct cut the two arguments are respectively $\log(Y/q)$ and $\log(Y/r)$. For the complementary cut they are $\log(bmrY/X)$ and $\log(bmqY/X)$. Thus each coordinate has its own one-dimensional knots. On each log-prime axis split at these knots, whose distinct spacings are $\gg 1/(bm)$, then subdivide each gap into intervals with log lengths between $Y^{-3\eta}$ and $2Y^{-3\eta}$, with the two infinite tails tiled likewise. Use whole rectangles lying inside $1/2 < t < 3/4$ and $q, r > Y^{1-2\eta}$. They cover all the nonzero integrand for $0.56 < t < 0.69$. In the interiors the weight is constant and neither open nor closed endpoints matter. Each relevant side at height Q has length $\gg QY^{-3\eta} \ge Q^{9/10}$, on fixing a sufficiently small η. The prime number theorem in short intervals gives $\gg |I|/\log X$ primes in each open physical interval I. Requiring unequal primes removes at most half their ordered pairs since both side counts tend uniformly to infinity. The change of variables in ordinary Lebesgue measure is $dq\, dr = X\, dt\, du/(bm)$ in absolute value, also when $T = tY$. Finally $z^K \gg_z z^{\omega(m)}$. This proves the lower bound

$$\frac{1}{(\log X)^2} \sum_m \frac{z^{\omega(m)}}{bm} \int_{0.56}^{0.69} \int_{\mathbb{R}} |J(u) + \mu(m)J(u - H)|^2 \, du\, dt$$

up to a fixed positive factor for the left side of (FS5).

Use Plancherel as in (MV4). For $1/\log b \le \xi \le 3/(2\log b)$, the mean over the indicated t interval of $|1 + \mu(m)e^{i\xi H}|^2$ is bounded below by a positive constant independent of X and m. Indeed $H = \log(Y^2/X) \pm \log t$, and the average of $t^{i\xi}$ has absolute value strictly less than one, uniformly on this compact frequency interval. The other factor $|(1 - e^{i\xi \log b})/\xi|^2$ is bounded below there. Average $\prod_{p|m} |1 - p^{i\xi}|^2$ over the same small cores as in (MV4), first allowing primes through $Y^{\eta'}$, then bounding the product m by Y^η by its expected logarithm. This gives $\gg_z l^{2z}$. In the present range $\log X \asymp l$, proving (FS5). The upper and shell estimates (FS2)–(FS4) already hold for $Y \le \sqrt{X}$, which proves the claimed full integer range. The independent prime sampling of q and r here uses the absence of a primality requirement on $X - n$.

11.5 A consequence above the second moment

For every fixed $s > 2$ one still has the lower bound

$$\mathbb{E}|M_R|^s \gg_s l^{2^{s/2}-1-s/2} \tag{FS6}$$

on all integer cutoffs and in the same fixed-radical or fixed-shift prime range. The implicit constant also depends on the fixed row parameters. Put $g = 2^{s/2}$, $z = g/2$. By (MV1) or (FS5) the tilted second moment is $\gg l^{2z-2}$; the multiplicative upper sieve gives $\mathbb{E}g^K \ll l^{g-1}$. Since $z^{s/(s-2)} = g$, Hölder gives $\mathbb{E}z^K M_R^2 \le (\mathbb{E}|M_R|^s)^{2/s}(\mathbb{E}g^K)^{1-2/s}$. Substitution gives (FS6) with a fixed positive constant. The same exponent, and the same Hölder step against $2^{(s/2)\omega(n)}$, appear in de la Bretèche–Tenenbaum's Theorem 1.5 (J. Lond. Math. Soc. 85 (2012), §4) as a lower bound for the moments of the localized supremum $\Delta(n,\mu)$ at every real order $s \ge 2$; there the tilted quadratic sum comes from an L^2 average over the position of the divisor interval. (FS6) is the fixed-cut version. For example the third moment has lower exponent $2\sqrt{2} - 5/2$ and the earlier upper exponent $2\sqrt{3} - 3$. GKM's fourth moment on integers already has exponent two in its asymptotic range; (FS6) alone gives exponent one there. The exact branch beyond two remains unresolved here.

References

[1] F. Dress, H. Iwaniec and G. Tenenbaum, Sur une somme liée à la fonction de Möbius, *J. Reine Angew. Math.* 340 (1983), 53–58. Author copy.

[2] R. R. Hall and G. Tenenbaum, The average orders of Hooley's Δ_r-functions, II, *Compositio Math.* 60 (1986), 163–186. Theorem 3.

[3] K. Ford, M. R. Gabdullin and K. Pratt, Integers with many divisors in a dyadic interval (2026). arXiv:2611.15534v1.

[4] A. Hildebrand, Über die punktweise Konvergenz von Ramanujan-Entwicklungen zahlentheoretischer Funktionen, *Acta Arith.* 44 (1984), 109–140. doi:10.4064/aa-44-2-109-140.

[5] K. Ford, Rough integers with a divisor in a given interval. arXiv:1901.02548.

[6] K. Ford, The distribution of integers with a divisor in a given interval, *Ann. of Math.* 168 (2008), 367–433. doi:10.4007/annals.2008.168.367.

[7] A. Granville, D. Koukoulopoulos and J. Maynard, Sieve weights and their smoothings, *Ann. Sci. Éc. Norm. Supér.* 54 (2021), 1089–1177. doi:10.24033/asens.2478.

[8] D. Koukoulopoulos, Divisors of shifted primes, *Int. Math. Res. Not.* 2010 (2010), 4585–4627. arXiv:0905.0163.

[9] R. de la Bretèche, F. Dress and G. Tenenbaum, Remarques sur une somme liée à la fonction de Möbius, *Mathematika* 66 (2020), 416–421. doi:10.1112/mtk.12021.

[10] R. de la Bretèche and G. Tenenbaum, Oscillations localisées sur les diviseurs, *J. Lond. Math. Soc.* 85 (2012), 669–693. doi:10.1112/jlms/jdr058.

[11] K. Henriot, Nair–Tenenbaum bounds uniform with respect to the discriminant, with erratum, New Theorem 5. doi:10.1017/S0305004114000280.

[12] P. Erdős and R. R. Hall, On the Möbius function, *J. Reine Angew. Math.* 315 (1980), 121–126.

[13] P. Letendre, Truncated convolution of Möbius function and multiplicative energy of an integer n, arXiv:1903.05629.

[14] O. Gorodetsky and L. Grimmelt, On an Erdős–Kac-type conjecture of Elliott, *Quart. J. Math.* 75 (2024), 695–713. arXiv:2307.13585.