hexagon:2611.00787v1

TWO SMOOTH SUMMANDS FOR EVERY LARGE INTEGER

WILL HOWE, GRANT HUFFMAN, JENSEN KOHLMEYER, AND LIAM KRUER

ABSTRACT. For every $\varepsilon > 0$, every sufficiently large integer N is a sum of two positive integers all of whose prime factors are at most N^{ε}. We prove this smooth-summand conjecture of Erdős using weighted dilation graphs, a uniform character-energy estimate and a positive marked principal term.

1. The smooth-summand theorem

Erdős asked whether every sufficiently large integer is a sum of two positive integers whose prime factors are smaller than any prescribed fixed power of that integer [4, p. 272]. We prove the following form of this conjecture.

Theorem 1.1. *For every $\varepsilon > 0$ there is an integer $N_0(\varepsilon)$ such that every integer $N \geq N_0(\varepsilon)$ admits positive integers a, b with*

$$N = a + b, \qquad p \mid ab \Longrightarrow p \leq N^{\varepsilon} \quad \text{for every prime } p.$$

For every fixed $\alpha > 0$, applying the theorem with $\varepsilon = 1/(2\alpha)$ gives the strict bound $p < N^{1/\alpha}$ once $N > 1$. This is the original formulation. The threshold formulation in [5, p. 55] and the uniform formulation in [6, Problem 59] are equivalent by absorbing the finitely many initial targets into the smoothness threshold.

The proof uses the uniform Dirichlet zero-free half-plane theorem of OpenAI [7] and the weighted-dilation methods developed in [8]. The reflected shift in the argument is the target $-N$. The prime-sum estimates, graph moment estimate and character-energy estimates below retain uniformity in this shift and in the corresponding moduli.

Throughout, a threshold may depend on fixed parameters chosen earlier in the argument. Every assertion of uniformity specifies the data allowed to vary after that threshold. We write $L = \log N$ when the target N is fixed.

2. Reduction to rough targets

For a positive integer $a > 1$, let $P^+(a)$ denote its largest prime factor, and put $P^+(1) = 1$. A positive integer is y-smooth if all its prime divisors are at most y, and is y-rough if all its prime divisors are greater than y. The integer 1 satisfies both conditions.

Lemma 2.1. *Let $y \geq 1$. Every y-smooth integer $N \geq 2$ is a sum of two positive y-smooth integers.*

Proof. Choose a prime divisor p of N and write $a = N/p$. Then

$$N = a + (p-1)a.$$

Date: November 13, 2026.

2020 *Mathematics Subject Classification.* 11N25, 11P32.

Key words and phrases. Smooth numbers, additive representations, Dirichlet L-functions, dilation graphs.

All authors contributed equally.

Electronic addresses. howew@purdue.edu (W. Howe), huffma97@purdue.edu (G. Huffman), jenwin@purdue.edu (J. Kohlmeyer), and kruerl@purdue.edu (L. Kruer).

Both summands are positive. Every prime divisor of a divides N, and every prime divisor of $p-1$ is at most $p-1<p\le y$. Thus both summands are y-smooth. □

Lemma 2.2. *Suppose that for every fixed $0<\eta<1/4$ there is $R_0(\eta)$ such that every $r\ge R_0(\eta)$ which is r^η-rough has a representation*

$$r=u+v,\qquad u,v\ge 1,\qquad P^+(uv)\le r^\eta.$$

Then, for every fixed $\varepsilon>0$, every sufficiently large integer N is the sum of two positive N^ε-smooth integers.

Proof. Fix $\varepsilon>0$ and set

$$\eta=\min\{\varepsilon/2,1/8\}.$$

Then $0<\eta<1/4$ and $\eta\le\varepsilon$. Take $N\ge 2$ sufficiently large that $N^\eta\ge R_0(\eta)$. Separate the prime factors of N at N^η, with multiplicity, by writing

$$d=\prod_{p\le N^\eta}p^{v_p(N)},\qquad r=N/d.$$

Here d is N^η-smooth, while every prime divisor of r exceeds N^η.

If $r=1$, then $N=d$ is N^η-smooth and Lemma 2.1 gives the required representation. If $r>1$, choose a prime divisor p of r. We have

$$r\ge p>N^\eta\ge R_0(\eta).$$

Moreover, $r\le N$, so each prime divisor of r exceeds $N^\eta\ge r^\eta$. The hypothesis therefore gives positive integers u,v with $r=u+v$ and $P^+(uv)\le r^\eta$. Consequently

$$N=du+dv.$$

Every prime divisor of either summand divides d, u, or v, and is therefore at most $N^\eta\le N^\varepsilon$. This proves the claim. □

3. Prime-sum estimates

Write $e(t)=\exp(2\pi it)$. Dirichlet characters are extended to all integers by zero on the nonunits of their modulus. The principal character modulo q is denoted by χ_0. Every prime sum below contains all primes in its displayed interval, with the lower endpoint excluded and the upper endpoint included. All absolute constants may be increased from one occurrence to the next.

3.1. The classical prime number theorem with a smooth weight.

Let Λ be the von Mangoldt function, and put

$$\psi(x)=\sum_{n\le x}\Lambda(n),\qquad \vartheta(x)=\sum_{p\le x}\log p.$$

The classical prime number theorem gives absolute constants $C,c>0$ such that, for sufficiently large x,

$$|\psi(x)-x|\le Cx\exp(-c\sqrt{\log x}). \tag{3.1}$$

We use this form of the theorem from [3]. The same estimate holds with ϑ in place of ψ. Indeed, Chebyshev's upper bound $\vartheta(t)\ll t$ gives

$$0\le\psi(x)-\vartheta(x)=\sum_{k\ge 2}\vartheta(x^{1/k})\ll\sqrt{x}.$$

Here the terms with $k\ge 3$ are $O(x^{1/3}\log(2x))$, which are $O(\sqrt{x})$, and $\sqrt{x}\le x\exp(-c\sqrt{\log x})$ once x is sufficiently large. Thus, after increasing C,

$$|\vartheta(x)-x|\le Cx\exp(-c\sqrt{\log x}). \tag{3.2}$$

Lemma 3.1 (Weighted prime replacement)**.** *There are absolute constants $C, c > 0$ and $a_0 \geq e$ such that the following holds. Let $a_0 \leq a \leq b$, let $W\colon [a,b] \to \mathbb{C}$ be continuously differentiable, and suppose $M \geq 0$ and $|W(t)| \leq M$ on $[a,b]$. Then*

$$\left| \sum_{a<p\leq b} W(p) - \int_a^b \frac{W(t)}{\log t}\, dt \right| \leq Cb\exp(-c\sqrt{\log a}) \cdot \left(2M + \int_a^b |W'(t)|\, dt + M\log(b/a) \right). \tag{3.3}$$

The threshold a_0 is independent of b and W.

Proof. Set $E(t) = \vartheta(t) - t$. Stieltjes integration by parts, with the right-continuous function ϑ, gives

$$\sum_{a<p\leq b} W(p) - \int_a^b \frac{W(t)}{\log t}\, dt = \frac{W(b)E(b)}{\log b} - \frac{W(a)E(a)}{\log a} - \int_a^b E(t) \left(\frac{W'(t)}{\log t} - \frac{W(t)}{t(\log t)^2} \right) dt.$$

By (3.2), throughout $[a,b]$ one has $|E(t)| \leq Cb\exp(-c\sqrt{\log a})$. Since $\log a \geq 1$, the two endpoint terms contribute at most $2M$ times this bound. The derivative terms contribute at most

$$\int_a^b |W'(t)|\, dt + M\int_a^b \frac{dt}{t} = \int_a^b |W'(t)|\, dt + M\log(b/a).$$

This proves (3.3), including $a = b$. □

3.2. **Uniform Dirichlet nonvanishing.** For $\operatorname{Re} s > 1$, use the normalization

$$L(s,\chi) = \sum_{n\geq 1} \frac{\chi(n)}{n^s}.$$

The following zero-free theorem is due to OpenAI.

Theorem 3.2 (OpenAI, [7], Theorem 1.1)**.** *For every positive modulus q and every Dirichlet character χ modulo q, the continued function $L(s,\chi)$ has no zero in $\operatorname{Re} s > 7/8$. For the principal character its pole at $s = 1$ is excluded from this assertion.*

For a nonprincipal character put

$$\mathcal{H}(q,\tau) = 1 + \log((q+2)(|\tau|+3)).$$

Lemma 3.3 (Logarithmic derivative)**.** *Uniformly over all positive moduli, nonprincipal characters, and complex s with $22/25 \leq \operatorname{Re} s \leq 2$,*

$$\left| \frac{L'}{L}(s,\chi) \right| \ll \mathcal{H}(q, \operatorname{Im} s). \tag{3.4}$$

Proof. The character sum over one complete period is zero. Hence $A_\chi(x) = \sum_{n\leq x} \chi(n)$ satisfies $|A_\chi(x)| \leq q+1$. Partial summation, followed by analytic continuation from $\operatorname{Re} s > 1$, gives

$$L(s,\chi) = s\int_1^\infty A_\chi(t) t^{-s-1}\, dt \qquad (\operatorname{Re} s > 0).$$

In particular,

$$|L(s,\chi)| \leq 2(q+1)|s| \qquad (\operatorname{Re} s \geq 1/2). \tag{3.5}$$

Fix $t \in \mathbb{R}$ and center a disk at $z_t = 2 + it$. Let

$$R_0 = \frac{9}{8} - \frac{2}{1600}, \qquad R_1 = \frac{9}{8} - \frac{3}{1600}, \qquad R_2 = \frac{9}{8} - \frac{8}{1600} = \frac{28}{25}.$$

The disk of radius R_0 lies in the half-plane in Theorem 3.2. It therefore has a holomorphic logarithm G of $L(s, \chi)$. Normalize it at z_t using the absolutely convergent Euler logarithm

$$G(z_t) = \sum_p \sum_{k \geq 1} \frac{\chi(p)^k}{k p^{k z_t}}.$$

Its absolute value is bounded by $\log \zeta(2)$, independently of q, t. On the disk of radius R_1, (3.5) gives

$$\operatorname{Re} G(s) = \log |L(s, \chi)| \leq \log\bigl(4(q+2)(|t|+3)\bigr).$$

Borel–Carathéodory on concentric intermediate disks and Cauchy's estimate for the derivative now give

$$|G'(s)| \ll 1 + \log\bigl((q+2)(|t|+3)\bigr) \qquad (|s - z_t| \leq R_2).$$

All radius gaps are positive absolute constants. The real segment $22/25 \leq \operatorname{Re} s \leq 2$, $\operatorname{Im} s = t$, lies in this last disk, and $G' = L'/L$. This proves the lemma. □

3.3. **Twisted prime sums.**

Lemma 3.4 (Smoothed and sharp twisted sums)**.** *Let $q \geq 1$, $\chi \neq \chi_0$ modulo q, and $\tau \in \mathbb{R}$. For $x \geq 1$,*

$$\left| \sum_{n \leq x} \Lambda(n) \chi(n) n^{i\tau} \left(1 - \frac{n}{x}\right) \right| \ll \mathcal{H}(q, \tau) x^{63/64}. \tag{3.6}$$

For $x \geq 2$,

$$\left| \sum_{p \leq x} \chi(p) p^{i\tau} \log p \right| \ll \mathcal{H}(q, \tau) x^{255/256}. \tag{3.7}$$

Both constants are absolute.

Proof. Write $a_n = \Lambda(n)\chi(n)n^{i\tau}$ and $T(x) = \sum_{n \leq x} a_n (1 - n/x)$. For $\operatorname{Re} s > 1$ the logarithmic derivative of the Euler product gives

$$\sum_{n \geq 1} \frac{a_n}{n^s} = -\frac{L'}{L}(s - i\tau, \chi).$$

The triangular Perron kernel gives the absolutely convergent identity

$$T(x) = \frac{1}{2\pi i} \int_{3/2 - i\infty}^{3/2 + i\infty} -\frac{L'}{L}(s - i\tau, \chi) \frac{x^s}{s(s+1)}\, ds. \tag{3.8}$$

Indeed, the inverse integral for a single term is $(1 - n/x)_+$. Its value at $n = x$ is zero. Absolute convergence of the Dirichlet series on the line $\operatorname{Re} s = 3/2$ justifies termwise integration.

Move the line in (3.8) to $\sigma_0 = 63/64$. Theorem 3.2 and $\chi \neq \chi_0$ make the integrand holomorphic throughout the intervening strip. Lemma 3.3 and

$$\mathcal{H}(q, t - \tau) \leq \mathcal{H}(q, \tau) + \log(1 + |t|)$$

give the uniform majorant

$$\left| -\frac{L'}{L}(\sigma + it - i\tau, \chi) \frac{x^{\sigma + it}}{(\sigma + it)(\sigma + 1 + it)} \right| \ll x^{3/2} \frac{\mathcal{H}(q, \tau) + \log(1 + |t|)}{1 + t^2} \quad (\sigma_0 \leq \sigma \leq 3/2).$$

The horizontal integrals tend to zero and both vertical integrals converge absolutely. On the new line, replace $x^{3/2}$ in this majorant by $x^{63/64}$. Its integral is $O(\mathcal{H}(q,\tau)x^{63/64})$, proving (3.6).

To remove the triangular weight, set $h = \lceil x^{127/128} \rceil$. The exact identity

$$\frac{(x+h)T(x+h) - xT(x)}{h} = \sum_{n \le x} a_n + \sum_{x < n \le x+h} a_n \frac{x+h-n}{h}$$

holds for every real $x \ge 1$. Since $|a_n| \le \log n$, $x^{127/128} \le h \le 2x^{127/128}$ and $x + h \le 3x$, (3.6) yields

$$\left| \sum_{n \le x} \Lambda(n)\chi(n)n^{i\tau} \right| \ll x^{127/128}(\mathcal{H}(q,\tau) + \log(3x)). \tag{3.9}$$

The discarded higher prime powers contribute at most

$$\sum_{\substack{p^k \le x \\ k \ge 2}} \log p \ll \sqrt{x}(\log(2x))^2.$$

Using $\log(3x) \ll x^{1/256}$ and $\mathcal{H}(q,\tau) \ge 1$ in (3.9) proves (3.7). □

Lemma 3.5 (Uniform reciprocal prime intervals)**.** *There is an absolute constant $C > 0$ such that, for every positive integer q, every nonprincipal character χ modulo q, every $\tau \in \mathbb{R}$, and all real $2 \le a \le b$,*

$$\left| \sum_{a < p \le b} \frac{\chi(p)p^{i\tau}}{p} \right| \le C\mathcal{H}(q,\tau)a^{-1/256}. \tag{3.10}$$

The upper endpoint b is unrestricted.

Proof. Let $A(t) = \sum_{p \le t} \chi(p)p^{i\tau} \log p$ and $v(t) = 1/(t \log t)$. Stieltjes integration by parts gives

$$\sum_{a < p \le b} \frac{\chi(p)p^{i\tau}}{p} = v(b)A(b) - v(a)A(a) - \int_a^b v'(t)A(t)\,dt.$$

By (3.7), the two endpoint terms are bounded by $C\mathcal{H}(q,\tau)a^{-1/256}$. For $t \ge 2$,

$$|v'(t)| = \frac{\log t + 1}{t^2(\log t)^2} \ll t^{-2}.$$

Thus the integral is bounded by

$$C\mathcal{H}(q,\tau) \int_a^\infty t^{-257/256}\,dt = 256C\mathcal{H}(q,\tau)a^{-1/256}.$$

The same right-continuous convention at both endpoints proves the assertion when an endpoint is a prime, as well as when $a = b$. □

Corollary 3.6 (The ranges used below)**.** *Fix $\kappa > 0$. For all sufficiently large L, uniformly for*

$$q \le e^{2L}, \qquad |\tau| \le e^{3L}, \qquad e^{L^\kappa} \le a \le b, \qquad \chi \ne \chi_0 \pmod q,$$

one has

$$\left| \sum_{a < p \le b} \frac{\chi(p)p^{i\tau}}{p} \right| \le \exp(-L^\kappa/512). \tag{3.11}$$

Moreover,

$$\left|\sum_{a<p\le b}\chi(p)p^{i\tau}\right| \le 2b\exp(-L^{\kappa}/512). \tag{3.12}$$

For fixed $D_1, D_2, A > 0$, the bound in (3.10) is at most L^{-A} when $q \le L^{D_1}$, $|\tau| \le L^{D_2}$ and $a \ge e^{L^{\kappa}}$, once the threshold in L is increased.

Proof. For $L \ge 1$ in the stated range, $\mathcal{H}(q,\tau) \le C(1+L)$. Hence

$$C\mathcal{H}(q,\tau)a^{-1/256} \le C'(1+L)\exp(-L^{\kappa}/256) \le \exp(-L^{\kappa}/512)$$

for sufficiently large L. This proves (3.11) uniformly in both endpoints. Put $B(t) = \sum_{a<p\le t}\chi(p)p^{i\tau}/p$ for $a \le t \le b$. A second partial summation gives

$$\sum_{a<p\le b}\chi(p)p^{i\tau} = bB(b) - \int_a^b B(t)\,dt.$$

The uniform bound for $B(t)$ proves (3.12). The last assertion follows because a fixed polynomial in L is smaller than $\exp(L^{\kappa}/512)$ for sufficiently large L. □

In particular, let $h > 0$, $R > 0$, $j \in \mathbb{Z}$, and $2 \le a_g \le b_g$. Freezing positive prime labels with product R in a logarithmic cell restricts the remaining prime to the interval

$$\left(\max\{a_g, e^{(j-1)h}/R\}, \min\{b_g, e^{jh}/R\}\right].$$

If this interval is empty, its sum is zero. Otherwise its lower endpoint is at least a_g, so Lemma 3.5 gives the bound with $a_g^{-1/256}$. Its constant and threshold are uniform in j and R. For two distinct primitive characters of conductor at most Q_0, their quotient on the least common multiple of the conductors is nonprincipal and has modulus at most Q_0^2. Thus the last assertion of Corollary 3.6 supplies the off-diagonal cell estimates when Q_0 is bounded by a fixed power of L. The large-modulus estimates (3.11)–(3.12) also give the uniform input for the reciprocal character energy in Section 6.

4. The invariant center and principal cancellation

Throughout this section, $0 < \varepsilon < 1/4$ and $0 < d_0 < 1/2$ are fixed. For a sufficiently large integer N put

$$\begin{aligned} &L = \log N, \qquad y = \exp(\varepsilon L), \qquad Q = \exp(L^{d_0}),\\ &D_Q = \prod_{p\le Q}(1-p^{-1}), \qquad D_* = \exp(L^{d_0}(\log L)^2). \end{aligned} \tag{4.1}$$

All accuracy exponents, frequency exponents and support constants in this section are fixed as N tends to infinity. Constants in the estimates may depend on these fixed parameters.

For $n \ge 1$ define its smooth and rough parts by

$$n^{\flat} = \prod_{p\le Q} p^{v_p(n)}, \qquad n^{\sharp} = \prod_{p>Q} p^{v_p(n)}, \qquad n = n^{\flat}n^{\sharp}.$$

Thus $n^{\flat}$ is Q-smooth and every prime divisor of $n^{\sharp}$ exceeds Q. We use the convention $P^+(1) = 1$. Write

$$h_0(n) = \mathbf{1}_{n=n^{\sharp}}, \qquad f_y(n) = \mathbf{1}_{P^+(n)\le y}, \qquad b_0(a) = \mu(a)\mathbf{1}_{P^+(a)\le Q}.$$

The function b_0 is supported on the divisors of the product of the primes at most Q. In particular,

$$h_0(n) = \sum_{a|n} b_0(a), \qquad D_Q = \sum_{a\ge 1} \frac{b_0(a)}{a}. \tag{4.2}$$

4.1. **Rough integers and reciprocal tails.**

Lemma 4.1. *For every fixed $A > 0$,*

$$\sum_{\substack{a>D_* \\ P^+(a)\le Q}} \frac{1}{a} \ll_A L^{-A}. \tag{4.3}$$

If $R_Q(x) = \sum_{1\le n\le x} h_0(n)$, then, for $x \ge 0$,

$$|R_Q(x) - D_Q x| \le \lceil D_*\rceil + 1 + x \sum_{\substack{a>D_* \\ P^+(a)\le Q}} \frac{1}{a}. \tag{4.4}$$

Consequently, for every fixed $\alpha > 0$ and $A > 0$,

$$\sup_{0\le x\le X} |R_Q(x) - D_Q x| \ll_{A,\alpha} XL^{-A} \qquad (X \ge e^{\alpha L}). \tag{4.5}$$

Furthermore,

$$H_Q(X) := \sum_{1\le m\le X} \frac{h_0(m)}{m} = D_Q \log X + O(L^{d_0}(\log L)^2) \qquad (1 \le X \le e^{2L}). \tag{4.6}$$

Proof. Set $\delta = 1/\log Q = L^{-d_0}$. Rankin's inequality gives

$$\sum_{\substack{a>D_* \\ P^+(a)\le Q}} \frac{1}{a} \le D_*^{-\delta} \prod_{p\le Q} (1 - p^{-1+\delta})^{-1}.$$

The shifted product is $O(\log Q)$. To see this, compare its logarithm with that of $\prod_{p\le Q}(1 - p^{-1})^{-1}$. The terms of degree one are bounded by

$$\sum_{p\le Q} \frac{p^\delta - 1}{p} \le e\delta \sum_{p\le Q} \frac{\log p}{p} \ll 1,$$

since $0 \le \delta \log p \le 1$. For large N, $\delta \le 1/4$, and the terms of degree at least two are bounded by a convergent series, uniformly in Q. Mertens' formula therefore proves the assertion about the product. As $D_*^{-\delta} = \exp(-(\log L)^2)$, this proves (4.3). We also obtain

$$D_Q^{-1} \asymp \log Q = L^{d_0}, \qquad \sum_{P^+(a)\le Q} \frac{1}{a} = D_Q^{-1}. \tag{4.7}$$

For an integer $j \ge 0$, (4.2) gives

$$R_Q(j) - D_Q j = \sum_{a\ge 1} b_0(a)(\lfloor j/a\rfloor - j/a).$$

For $a \le D_*$ bound each rounding error by one. There are at most $\lceil D_*\rceil$ such positive integers. For $a > D_*$ use $|\lfloor j/a\rfloor - j/a| \le j/a$, and then enlarge the absolute coefficient sum to the sum over all Q-smooth integers. Replacing j by $\lfloor x\rfloor$ costs at most $D_Q \le 1$. This proves (4.4). Since $\log D_* = o(L)$, the quantity $D_* + 2$ is smaller than $e^{\alpha L}L^{-A}$ for every fixed $\alpha, A > 0$ and sufficiently large N. Increasing the accuracy in (4.3) proves (4.5).

Partial summation gives, for $X \geq 1$,

$$H_Q(X) = \frac{R_Q(X)}{X} + \int_1^X \frac{R_Q(t)}{t^2}\, dt.$$

Here $0 \leq R_Q(t) \leq t$ and $0 < D_Q \leq 1$, so $|R_Q(t) - D_Q t| \leq t$. On $1 \leq t \leq D_*$ this contributes at most $\log D_*$ to the integral of the discrepancy divided by t^2. For $D_* \leq t \leq X \leq e^{2L}$, use (4.4) with the tail bounded by L^{-3}. The two resulting contributions are $O(1)$ and $O(L^{-3}\log X) = O(1)$. The endpoint $R_Q(X)/X$ is at most one. The same argument with an empty second range covers $X < D_*$ and proves (4.6). □

4.2. **The density and its point values.** Let $h(t) = 1/t$ for $t \geq 1$ and $h(t) = 0$ otherwise. Define $A_1 = h$ and, recursively,

$$A_{k+1}(u) = \int_{\mathbb{R}} A_k(u-t)h(t)\, dt.$$

Each such integral is over a bounded interval when u is fixed. Thus $A_k = h^{*k}$, with the right-continuous value at the jump of A_1. Define also

$$I_k(u) = \int_{\substack{t_1,\dots,t_k \geq 1 \\ t_1+\cdots+t_k \leq u}} \frac{dt_1 \cdots dt_k}{t_1 \cdots t_k}, \qquad I_0(u) = \mathbf{1}_{[0,\infty)}(u).$$

Then $I_k(u) = \int_{-\infty}^u A_k(t)\, dt$ for $k \geq 1$. In particular,

$$A_1(u) = \begin{cases} 0, & u < 1, \\ 1/u, & u \geq 1. \end{cases} \tag{4.8}$$

For $k \geq 2$, eliminating the last variable gives the pointwise formula

$$A_k(u) = \int_{\substack{t_1,\dots,t_{k-1} \geq 1 \\ t_1+\cdots+t_{k-1} \leq u-1}} \frac{dt_1 \cdots dt_{k-1}}{t_1 \cdots t_{k-1}(u - t_1 - \cdots - t_{k-1})}.$$

The convolution identity and symmetry of these coordinates give

$$uA_k(u) = kI_{k-1}(u-1) \qquad (u > 0). \tag{4.9}$$

Indeed, set $t_k = u - \sum_{j<k} t_j$ in the preceding integral and multiply by $u = \sum_j t_j$. Cancelling the factor t_k gives $I_{k-1}(u-1)$. Permuting the k variables preserves the integration domain and the coordinate measure, so the other $k-1$ summands have the same value. Formula (4.8) gives the case $k = 1$, including its endpoint value. The functions A_k vanish below k. They are bounded and of bounded variation on every fixed interval $(-\infty, U]$. Indeed, (4.9) expresses them, on $u \geq 1$, as the product of the bounded increasing function $kI_{k-1}(u-1)$ and $1/u$. The convolutions A_k are continuous for $k \geq 2$. On any compact u-interval one may truncate h to a bounded, integrable function. Continuity then follows from continuity of translations in L^1 and the boundedness of the other convolution factor.

For $x > 0$ define

$$g_0(x) = D_Q + \sum_{k \geq 1} \frac{(-1)^k}{k! \log y} \sum_{m \geq 1} \frac{h_0(m)}{m} A_k\left(\frac{\log(x/m)}{\log y}\right). \tag{4.10}$$

Both sums are locally finite. A nonzero summand requires $y^k m \leq x$. Thus g_0 is right-continuous at every positive point, with the value in (4.8) retained when $x = my$.

Lemma 4.2. *Uniformly on $e^{L/2} \leq x \leq e^{2L}$,*

$$\frac{g_0(x)}{D_Q} = \rho\left(\frac{\log x}{\log y}\right) + O_{\varepsilon,d_0}(L^{2d_0-1}(\log L)^2), \tag{4.11}$$

where ρ is the Dickman function. Moreover,

$$\operatorname{Var}_{[e^{L/2},e^{2L}]}(g_0/D_Q) \ll_{\varepsilon,d_0} 1. \tag{4.12}$$

For sufficiently large N, $g_0(x)/D_Q$ lies in $[0,1]$ throughout this interval.

Proof. Put $\ell = \log y$ and $U = 2/\varepsilon$, and choose a fixed integer $K > U$. The bounded-variation norms of $A_1, \dots, A_K$ on $(-\infty, U]$ are constants depending only on ε. Write $E_Q = CL^{d_0}(\log L)^2$, increasing C so that $E_Q \ge 1$ and the error in (4.6) is at most E_Q.

Apply Stieltjes partial summation to the harmonic cumulative function $H_Q(t)$ and the kernel $A_k(u - \log t/\ell)$, where $u = \log x/\ell \le U$. Its support in $t \ge 1$ is bounded by x, and its supremum and variation are bounded in terms of ε. Replacing $H_Q(t)$ by $D_Q \log t$ therefore gives

$$\sum_{m\ge 1} \frac{h_0(m)}{m} A_k\Big(u - \frac{\log m}{\ell}\Big) = D_Q \int_1^\infty A_k\Big(u - \frac{\log t}{\ell}\Big)\frac{dt}{t} + O_\varepsilon(E_Q) = D_Q \ell I_k(u) + O_\varepsilon(E_Q). \tag{4.13}$$

For continuous bounded-variation kernels the error bound follows directly by integrating the discrepancy in (4.6). For $k = 1$, put $Y = x/y$ and $s(t) = (u - \log t/\ell)^{-1}$ on $[1, Y]$. The sum terminates at Y, with that endpoint included. Writing $E(t) = H_Q(t) - D_Q \log t$, ordinary partial summation gives its error as

$$E(Y)s(Y) - \int_1^Y E(t)s'(t)\,dt.$$

Here $s(Y) = 1$, s is increasing, and $\int_1^Y s'(t)\,dt \le 1$, so the error is at most $2E_Q$. This calculation retains the displayed value of A_1 at an integer endpoint. In the range under consideration $Y > 1$.

Substitution into (4.10) gives

$$\frac{g_0(x)}{D_Q} = \sum_{k\ge 0} \frac{(-1)^k}{k!} I_k(u) + O_\varepsilon\Big(\frac{E_Q}{D_Q \ell}\Big).$$

Denote the finite sum by $\sigma(u)$. It is one on $[0,1]$, and (4.9) gives

$$u\sigma'(u) = \sum_{k\ge 1} \frac{(-1)^k}{(k-1)!} I_{k-1}(u-1) = -\sigma(u-1) \qquad (u > 1).$$

Successive integration on the intervals $[j, j+1]$ identifies σ with the Dickman function ρ. Together with (4.7), this proves (4.11).

For the variation assertion, use the finite expression in (4.10) on the whole interval under consideration. The variation of each composition of A_k with $\log(x/m)/\ell$ is at most the variation of A_k on $(-\infty, U]$. All coefficients in the absolute variation bound are nonnegative. Consequently,

$$\operatorname{Var}_{[e^{L/2},e^{2L}]}(g_0/D_Q) \le \frac{H_Q(e^{2L})}{D_Q\ell} \sum_{k=1}^K \frac{\operatorname{Var}_{(-\infty,U]} A_k}{k!} \ll_\varepsilon 1 + \frac{E_Q}{D_Q\ell} \ll_{\varepsilon,d_0} 1.$$

Finally, $u \in [1/(2\varepsilon), 2/\varepsilon]$, a fixed compact subinterval of $(1,\infty)$. The Dickman function is continuous, positive and strictly below one on this interval. Its delay equation gives $u\rho(u) = \int_{u-1}^u \rho(t)\,dt$ for $u > 1$. The two sides agree at $u = 1$ and have equal derivatives for $u > 1$. A first zero would contradict this identity, since $\rho = 1$ on $[0,1]$ and the integral at that zero would be positive. The delay equation then shows that ρ is strictly decreasing on $(1,\infty)$. Since $2d_0 - 1 < 0$, the error in (4.11) tends to zero. The final assertion follows. □

Define the bounded extension

$$\mathcal{M}(x) = \begin{cases} g_0(x)/D_Q, & e^{L/2} \le x \le e^{2L}, \\ \rho(1/\varepsilon), & \text{otherwise}, \end{cases} \qquad F(n) = f_y(n^\sharp) - \mathcal{M}(n^\sharp) \quad (n \ge 1). \tag{4.14}$$

For large N, $|F(n)| \le 1$. Unique smooth–rough factorization gives

$$F(an) = F(n) \qquad (a \text{ is } Q\text{-smooth},\ n \ge 1). \tag{4.15}$$

Also $f_y(n^\sharp) = f_y(n)$ once $Q \le y$.

4.3. **Right-continuous Stieltjes comparison.**

Lemma 4.3. *Let $X > 0$, and let $v\colon \mathbb{R} \to \mathbb{C}$ be right-continuous, of bounded variation, and compactly supported in $(0, X)$. Then*

$$\left| \sum_{1 \le n \le X} h_0(n) v(n) - D_Q \int_0^X v(x)\, dx \right| \le \Delta_Q(X)\big(2\|v\|_\infty + \operatorname{Var}_{\mathbb{R}} v\big), \tag{4.16}$$

$$\Delta_Q(X) = \sup_{0 \le x \le X} |R_Q(x) - D_Q x|.$$

The sum uses the specified right-continuous values of v at all atoms.

Proof. Define

$$R_{Q,X}(t) = \sum_{1 \le n \le X} h_0(n) \mathbf{1}_{n \le t}, \qquad c_X(t) = \max(0, \min(X, t)),$$

$$E(t) = R_{Q,X}(t) - D_Q c_X(t).$$

The function E is right-continuous and of bounded variation on $\mathbb{R}$. It vanishes below zero, is constant above X, and satisfies $\sup_{t \in \mathbb{R}} |E(t)| \le \Delta_Q(X)$. Its Stieltjes measure is

$$dE = \sum_{1 \le n \le X} h_0(n) \delta_n - D_Q \mathbf{1}_{(0,X)}(t)\, dt.$$

For two right-continuous bounded-variation functions, the product rule on a closed interval gives

$$\int_{[a,b]} v(t)\, dE(t) = v(b)E(b) - v(a-)E(a-) - \int_{[a,b]} E(t-)\, dv(t). \tag{4.17}$$

At a jump, the identity is $\Delta(vE) = v\,\Delta E + E_-\,\Delta v$, which verifies the one-sided conventions in (4.17). Both endpoint terms are included. The left limits of E have modulus at most $\Delta_Q(X)$ as well. Taking absolute values in (4.17) proves (4.16). For the stated support one can choose endpoints where both boundary values of v vanish. The atomic integral is exactly the sum of $h_0(n)v(n)$, including every threshold atom. □

Fix a support constant $C \ge 1$. We use complex C^1 test functions on $\mathbb{R}$ satisfying

$$\operatorname{supp} w \subset [Z/C, CZ], \qquad \|w\|_\infty \le L^B, \qquad \int_{\mathbb{R}} |w'(x)|\, dx \le L^B, \tag{4.18}$$

where $B \ge 0$ is fixed. In particular $\operatorname{Var}_{\mathbb{R}} w \le L^B$. All such tests are allowed to depend on N and Z.

Lemma 4.4. *For every fixed $A > 0$ and $B \ge 0$, uniformly over (4.18) and*

$$e^{5L/8} \le Z \le e^{11L/8}, \tag{4.19}$$

one has

$$\sum_{n \ge 1} h_0(n) \mathcal{M}(n) w(n) = \int_0^\infty g_0(x) w(x)\, dx + O_{A,B,C,\varepsilon,d_0}(ZL^{-A}). \tag{4.20}$$

Proof. For large N, $[Z/C, CZ]$ lies strictly inside $[e^{L/2}, e^{2L}]$. Extend g_0/D_Q from this latter interval by its left endpoint value on the left and its right endpoint value on the right. Call the resulting function $\widetilde{\mathcal{M}}$. It is right-continuous, bounded by one in modulus, and has uniformly bounded total variation by Lemma 4.2. The function $v = \widetilde{\mathcal{M}}w$ is right-continuous and satisfies

$$\|v\|_\infty \le L^B, \qquad \operatorname{Var}_{\mathbb{R}} v \le \|w\|_\infty \operatorname{Var}_{\mathbb{R}} \widetilde{\mathcal{M}} + \|\widetilde{\mathcal{M}}\|_\infty \operatorname{Var}_{\mathbb{R}} w \ll_{\varepsilon,d_0} L^B.$$

It agrees with $\mathcal{M}w$ everywhere. Take $X = \lceil 2CZ \rceil$. Then $X \asymp_C Z$, the support of v lies inside $(0, X)$, and $X \ge e^{L/2}$ for large N. Apply Lemma 4.3 and (4.5), choosing its accuracy after A and B. On the support of w the identity $D_Q\mathcal{M} = g_0$ is exact, giving (4.20). □

4.4. **The uncentered prime expansion.**

Lemma 4.5. *Under* (4.18) *and* (4.19)*, for every fixed* $A > 0$*,*

$$\sum_{n\ge1} h_0(n) f_y(n) w(n) = \int_0^\infty g_0(x) w(x)\, dx + O_{A,B,C,\varepsilon,d_0}(ZL^{-A}). \tag{4.21}$$

Proof. Put $X = \lceil 2CZ \rceil$ and fix an integer $K > 2/\varepsilon$. For large N, $y \le X \le e^{2L}$ and $Q \le y$. An integer at most X has at most K distinct prime divisors greater than y. Finite inclusion–exclusion therefore gives

$$\sum_n h_0(n) f_y(n) w(n) = \sum_{k=0}^{K} \frac{(-1)^k}{k!} \sum_{\substack{y<p_1,\dots,p_k\le X \\ p_1,\dots,p_k \text{ distinct}}} \sum_{m\ge1} h_0(m) w(mp_1\cdots p_k). \tag{4.22}$$

Removing a selected prime exceeding $y \ge Q$ preserves the roughness condition. The residual integer m in this formula may itself be divisible by a selected prime.

Write $H = 1 + \log X \ll L$. For $k \ge 2$, allowing repeated prime coordinates in (4.22) introduces an error bounded by

$$\begin{aligned} L^B \binom{k}{2} X \left(\sum_{y<p\le X} \frac{1}{p^2}\right) \left(\sum_{y<p\le X} \frac{1}{p}\right)^{k-2} &\le L^B \binom{k}{2} \frac{X}{y} H^{k-1} \\ &\ll_{\varepsilon,A,B} XL^{-A}. \end{aligned}$$

Indeed, after fixing the prime coordinates the number of possible positive m is at most $X/(p_1\cdots p_k)$. The union over the chosen repeated pair covers every tuple with a repetition, and the factor $1/k! \le 1$ has been discarded in this upper bound. The last inequality follows from $y = e^{\varepsilon L}$ and the fixed bound $k \le K$.

For the prime replacement, let

$$\nu_{\mathrm{d}} = \sum_{y<p\le X} \delta_p, \qquad d\nu_{\mathrm{c}}(t) = \mathbf{1}_{(y,X]}(t) \frac{dt}{\log t}.$$

The classical prime number theorem with its exponential remainder, followed by partial summation, gives, for $2 \le a \le b$ and a complex C^1 function W,

$$\left| \sum_{a<p\le b} W(p) - \int_a^b \frac{W(t)}{\log t}\, dt \right| \ll be^{-c\sqrt{\log a}} \left(2\|W\|_\infty + \int_a^b |W'(t)|\, dt + \|W\|_\infty \log(b/a) \right). \tag{4.23}$$

Here the supremum is taken on $[a, b]$, and $c > 0$ is absolute. To obtain this form, write $E_\pi(t) = \pi(t) - \int_2^t ds/\log s$. The prime number theorem gives $E_\pi(t) \ll t\exp(-c\sqrt{\log t})$ for

$t \geq 2$, and partial summation gives the exact identity

$$\sum_{a<p\leq b} W(p) - \int_a^b \frac{W(t)}{\log t}\,dt = W(b)E_\pi(b) - W(a)E_\pi(a) - \int_a^b E_\pi(t)W'(t)\,dt.$$

Since $|E_\pi(t)| \ll b\exp(-c\sqrt{\log a})$ on $[a,b]$, this proves (4.23). The prime number theorem used here is given in [3].

Expand $\nu_{\mathrm{d}}^{\otimes k} - \nu_{\mathrm{c}}^{\otimes k}$ as the sum of k products containing one factor $\nu_{\mathrm{d}} - \nu_{\mathrm{c}}$, with the preceding factors discrete and the following factors continuous. Fix m and all prime coordinates except that distinguished factor, and put $U = m\prod_{j\neq i} p_j$. Its test is the original function $W(t) = w(Ut)$. The support restricts t to $t \leq X/U$. If $X/U \leq y$, this term is zero. Otherwise (4.23), with $a = y$ and $b = X/U$, applies and

$$\|W\|_\infty \leq L^B, \qquad \int_y^{X/U} |W'(t)|\,dt \leq \int_{\mathbb{R}} |w'(x)|\,dx \leq L^B.$$

Its error is at most

$$\frac{X}{U}L^{B+1}\exp(-c_\varepsilon\sqrt{L}).$$

Both remaining prime measures have reciprocal mass at most H. Summing m^{-1} for $m \leq X$ costs at most H as well. Consequently the total replacement error, summed over all the fixed orders $k \leq K$, is

$$\ll_{\varepsilon,C} XL^{B+K+2}\exp(-c_\varepsilon\sqrt{L}) \ll_{A,B,C,\varepsilon} XL^{-A}.$$

Every application of (4.23) is made to $w(Ut)$ before integrating the other coordinates. Fubini's theorem applies because the coordinates are restricted to bounded intervals and the integrands are absolutely integrable.

It remains to identify the continuous density. The support of w and $m \geq 1$ imply that every contributing prime variable is at most $CZ < X$. Thus the upper truncation at X may be removed from the continuous integrals. With $\ell = \log y$, set $p_j = e^{\ell t_j}$. Then

$$\prod_{j=1}^k \frac{dp_j}{\log p_j} = e^{\ell(t_1+\cdots+t_k)}\prod_{j=1}^k \frac{dt_j}{t_j}.$$

The pushforward of $\prod_j \mathbf{1}_{t_j\geq 1}\,dt_j/t_j$ under addition has density A_k. Next put $x = me^{\ell u}$. The two substitutions give

$$\int_{p_1,\ldots,p_k>y} w(mp_1\cdots p_k)\prod_{j=1}^k \frac{dp_j}{\log p_j} = \frac{1}{m\ell}\int_0^\infty A_k\Big(\frac{\log(x/m)}{\ell}\Big)\,w(x)\,dx. \tag{4.24}$$

The boundary hyperplanes have zero Lebesgue measure. The factor $1/k!$ continues to come from the ordered inclusion–exclusion expansion in (4.22). Every contributing residual satisfies $m \leq X$. All k and m sums meeting the compact support are therefore finite, so they may be interchanged with the final integral. The term $k = 0$ is $D_Q\int w + O_A(XL^{-A})$ by Lemma 4.3, applied to $v = w$. Combining these facts with (4.10) proves (4.21). □

Subtracting Lemma 4.4 from Lemma 4.5 gives the following rough-integer estimate.

Corollary 4.6. *For every fixed $A > 0$, uniformly under (4.18) and (4.19),*

$$\sum_{n\geq 1} h_0(n)F(n)w(n) \ll_{A,B,C,\varepsilon,d_0} ZL^{-A}. \tag{4.25}$$

4.5. **Arbitrary smooth-part coefficients.**

Theorem 4.7. *Let $A > 0$, $B \ge 0$ and $C \ge 1$ be fixed. Suppose that b is any complex function on the positive Q-smooth integers with $|b(a)| \le L^B$. Uniformly for tests satisfying (4.18) and scales*

$$e^{3L/4} \le Z \le e^{5L/4}, \tag{4.26}$$

one has

$$\sum_{n\ge1} F(n)b(n^\flat)w(n) \ll_{A,B,C,\varepsilon,d_0} ZL^{-A}. \tag{4.27}$$

The function b and the test w may depend arbitrarily on N, Z and additional parameters subject to the stated uniform bounds.

Proof. Unique smooth–rough factorization and (4.15) give the exact identity

$$\sum_{n\ge1} F(n)b(n^\flat)w(n) = \sum_{\substack{a\ge1\\ P^+(a)\le Q}} b(a) \sum_{m\ge1} h_0(m)F(m)w(am). \tag{4.28}$$

For $a \le D_*$, the test $w_a(t) = w(at)$ has support in $[(Z/a)/C, C(Z/a)]$, and

$$\|w_a\|_\infty \le L^B, \qquad \int_{\mathbb{R}} |w_a'(t)|\,dt = \int_{\mathbb{R}} |w'(x)|\,dx \le L^B.$$

Since $\log D_* = o(L)$, for large N the outer band (4.26) and $1 \le a \le D_*$ imply

$$e^{5L/8} \le Z/a \le e^{11L/8}.$$

Apply Corollary 4.6 with accuracy $A_0 > A + B + d_0 + 1$, retaining the same Q, y, L and center. The contribution of $a \le D_*$ to (4.28) is at most

$$\ll ZL^{B-A_0} \sum_{P^+(a)\le Q} \frac{1}{a} \ll ZL^{B+d_0-A_0} \ll ZL^{-A}.$$

For $a > D_*$, use $|F| \le 1$, $|w| \le L^B$, and the exact bound $\#\{m \ge 1 : am \le CZ\} = \lfloor CZ/a \rfloor \le CZ/a$. The remaining contribution is at most

$$CZL^{2B} \sum_{\substack{a>D_*\\ P^+(a)\le Q}} \frac{1}{a} \ll_A ZL^{-A}$$

by (4.3), with its accuracy chosen after A and B. This proves (4.27). □

4.6. **The external target-unit restriction.**

Lemma 4.8. *Let M and N be positive integers, and suppose that every prime divisor of N exceeds $P > 0$. Then*

$$\#\{1 \le n \le M : (n, N) > 1\} \le \sum_{p\mid N} \lfloor M/p \rfloor \le \omega(N)M/P, \tag{4.29}$$

where $\omega(N)$ counts distinct prime factors. For a finite set $S \subset [a, b] \cap \mathbb{N}$, with $0 \le a \le b$, and $|c_n| \le H$ on S, one also has

$$\sum_{\substack{n\in S\\ (n,N)>1}} |c_n|^2 \le H^2\omega(N)\left(\frac{b-a}{P} + 1\right). \tag{4.30}$$

The latter estimate also holds when every prime divisor of N is at least P.

Proof. Every nonunit has a common prime divisor with N. The union bound over these distinct primes gives the first inequality in (4.29). The count for one prime is exactly $\lfloor M/p \rfloor$. For the interval version, any two consecutive multiples of p are separated by at least p. If there are j such multiples in $[a, b]$, then $(j-1)p \le b-a$. Summing this count over $p \mid N$ and multiplying by H^2 proves (4.30). □

Theorem 4.9. *In addition to the hypotheses of Theorem 4.7, fix $\eta > 0$ and $D \ge 0$, and let r be a positive integer such that*

$$P^-(N) > N^\eta, \qquad 1 \le r \le L^D.$$

Here $P^-(N)$ denotes the least prime factor of N. Then, for every fixed $A > 0$,

$$\sum_{\substack{n \ge 1 \\ (n, Nr)=1}} F(n) b(n^\flat) w(n) \ll_{A,B,C,D,\varepsilon,d_0,\eta} ZL^{-A}, \tag{4.31}$$

uniformly in r, b, w and the outer scale band (4.26).

Proof. For large N, $L^D \le Q$, so r is Q-smooth and $(r, n^\sharp) = 1$. Therefore

$$(n, r) = 1 \quad \Longleftrightarrow \quad (n^\flat, r) = 1.$$

The coefficient $b_r(a) = b(a)\mathbf{1}_{(a,r)=1}$ has the same bound as b. Theorem 4.7 applied to b_r estimates the sum restricted by $(n, r) = 1$.

The hypothesis on the least prime factor gives $\omega(N) < 1/\eta$, by comparing the product of the distinct prime divisors with N. Take $M = \lceil 2CZ \rceil \le 3CZ$, which contains the support of the sum. Lemma 4.8 shows that imposing $(n, N) = 1$ changes this sum in modulus by at most

$$L^{2B} \#\{1 \le n \le M : (n, N) > 1\} \le \frac{3C}{\eta} ZL^{2B} N^{-\eta} \ll_{A,B,C,\eta} ZL^{-A}.$$

Increasing the accuracy used for the first sum proves (4.31). □

Corollary 4.10. *Let $\psi \in C_c^1((0,\infty); \mathbb{C})$ and $\eta > 0$ be fixed. For fixed nonnegative B, E, D and every fixed $A > 0$, suppose that $P^-(N) > N^\eta$, that r is a positive integer with $r \le L^D$, and that Z lies in (4.26). Then*

$$\sum_{\substack{n \ge 1 \\ (n, Nr)=1}} F(n) b(n^\flat) \psi(n/Z) e(\beta n)\, n^{i\zeta} \ll_{A,B,E,D,\psi,\varepsilon,d_0,\eta} ZL^{-A},$$

uniformly for $|b(a)| \le L^E$ and real β, ζ satisfying $|\beta| \le L^B/Z$ and $|\zeta| \le L^B$.

Proof. Choose a fixed $C \ge 1$ such that $\operatorname{supp} \psi \subset [1/C, C]$. Put $w(x) = \psi(x/Z) e(\beta x)\, x^{i\zeta}$ for $x > 0$ and $w(x) = 0$ for $x \le 0$. The cutoff vanishes near zero, so this function is C^1 on $\mathbb{R}$. Its supremum is at most $\|\psi\|_\infty$, and differentiation on its positive support gives

$$\int_{\mathbb{R}} |w'(x)|\, dx \le \int_0^\infty |\psi'(t)|\, dt + 2\pi|\beta| Z \int_0^\infty |\psi(t)|\, dt + |\zeta| \int_0^\infty \frac{|\psi(t)|}{t}\, dt.$$

Thus both test bounds are at most L^{B+1} for sufficiently large N. Use Theorem 4.9 with common exponent $\max(E, B+1)$. □

Corollary 4.11. *Fix $\eta > 0$. Under $P^-(N) > N^\eta$, for sufficiently large N and every positive integer r,*

$$\frac{N}{\varphi(N)} \le 2, \qquad \frac{Nr}{\varphi(Nr)} \le 2r.$$

In particular the second ratio is at most $2L^D$ for $r \le L^D$.

Proof. Applying (4.29) with $M = N$ gives $N - \varphi(N) \le \eta^{-1} N^{1-\eta}$, and hence $\varphi(N) \ge N/2$ for large N. The Euler product gives

$$\frac{Nr}{\varphi(Nr)} \le \frac{N}{\varphi(N)} \frac{r}{\varphi(r)} \le 2r,$$

since $\varphi(r) \ge 1$. □

5. Constructed dilation graphs and their moments

The weighted-dilation and residual-kernel methods used here come from [8]. In this section

$$q = \frac{1}{256}, \qquad \ell = 1, \qquad (a,b,c,d) = \left(\frac{27}{100}, \frac{36}{100}, \frac{39}{100}, \frac{46}{100}\right).$$

The parameter L tends to infinity. Write $T = \log L$. For $\alpha < \beta$, put

$$s_{\alpha,\beta} = \lfloor (\beta - \alpha)T \rfloor, \qquad x_{\alpha,j} = \exp(\exp(\alpha T + j)),$$
$$\mathcal{P}_{\alpha,j} = (x_{\alpha,j}, x_{\alpha,j+1}] \cap \mathbb{P} \quad (0 \le j < s_{\alpha,\beta}).$$

The small groups $\mathcal{S}$ are the groups with $(\alpha,\beta) = (a,b)$. The big groups $\mathcal{B}$ are those with $(\alpha,\beta) = (c,d)$. Set

$$s_S = |\mathcal{S}|, \quad s_B = |\mathcal{B}|, \quad s = s_S + s_B, \qquad V_g = \sum_{p \in \mathcal{P}_g} \frac{1}{p}, \qquad \mu_g(p) = \frac{1}{V_g p}.$$

The prime number theorem, or the corresponding form of Mertens' theorem, gives, for sufficiently large L,

$$\frac{1}{2} \le V_g \le \frac{3}{2}, \qquad \frac{9}{200}T \le s_S \le \frac{9}{100}T, \qquad \frac{7}{200}T \le s_B \le \frac{7}{100}T. \tag{5.1}$$

Every small prime is at most $\exp(L^{36/100})$, and every big prime exceeds

$$P_* = \exp(L^{39/100}).$$

All prime intervals have the displayed endpoint convention. Thus their prime sets are pairwise disjoint.

5.1. **The two row operators.** Fix $J \ge 1$ and put $M = J+1$. Slots are numbered $1, \dots, M$. A pattern ν specifies a set $C_{\nu,g} \subseteq \{1, \dots, J\}$. Its shared and unshared slots are

$$S_{\nu,g} = \{1, \dots, J\} \setminus C_{\nu,g}, \qquad U_{\nu,g} = \{1, \dots, M\} \setminus S_{\nu,g}, \qquad t_{\nu,g} = |U_{\nu,g}| = 1 + |C_{\nu,g}|.$$

We always require $C_{\nu,g} = \varnothing$ for $g \in \mathcal{S}$. For an array $\mathbf{p} = (p_{g,i})$ and a free array $\mathbf{r} = (r_{g,i})_{i \in C_{\nu,g}}$, define

$$D_\nu(\mathbf{p}, \mathbf{r}) = \prod_g \left(\prod_{i \in S_{\nu,g}} p_{g,i} \prod_{i \in C_{\nu,g}} r_{g,i} \right). \tag{5.2}$$

Its big-group part is denoted by $D_{\nu,B}$. Repetition is allowed in every free array.

The ideal Hilbert space is

$$\mathcal{H}_B = L^2\left(\prod_{g \in \mathcal{B}} \mathcal{P}_g^M, \bigotimes_{g \in \mathcal{B}} \mu_g^{\otimes M} \right).$$

For fixed source labels, the ideal row operation sums over the target labels in $U_{\nu,g}$ and the free labels in $C_{\nu,g}$, using their product laws μ_g. Shared target labels equal the source labels. The summand is multiplied by

$$c_\nu(\mathbf{p}_B, \mathbf{p}'_B, \mathbf{r}_B) D_{\nu,B}^{i\zeta} e(\Theta D_{\nu,B}).$$

Let S average all slot permutations independently in each big group. It is the orthogonal projection onto the symmetric vectors of $\mathcal{H}_B$. The complete ideal operator is the sum of the preceding row operations, with S applied at both ends. Denote it by $\mathcal{T}_\zeta(\Theta)$.

For a finite set $\mathcal{W} \subset \mathbb{Z}$, the physical vertices are

$$\mathcal{V}_\mathcal{W} = \{(z, \mathbf{p}) \colon z \in \mathcal{W}, \quad p_{g,i} \in \mathcal{P}_g, \quad p_{g,i} \mid z, \quad p_{g,1}, \dots, p_{g,M} \text{ are distinct for every } g\}.$$

Every vertex has mass

$$\tau = \prod_g V_g^{-M}.$$

The physical Hilbert space is $L^2(\mathcal{V}_\mathcal{W}, \tau)$. For $x = (z, \mathbf{p})$, set

$$d_q(x) = q^{(\omega_\mathcal{P}(z) - Ms)/2}.$$

Here $\omega_\mathcal{P}(z) = \sum_g \sum_{p \in \mathcal{P}_g} \mathbf{1}_{p|z}$. The vertex conditions imply $\omega_\mathcal{P}(z) \geq Ms$.

For $y = (z', \mathbf{p}')$, the raw row kernel of pattern ν is

$$K_{\nu,k}^{\zeta,\theta}(x, y) = d_q(x) d_q(y) \prod_g V_g^{-t_{\nu,g}} \sum_{\mathbf{r}} \left(\prod_g \prod_{i \in C_{\nu,g}} \mu_g(r_{g,i}) \right)$$
$$\times c_\nu(\mathbf{p}_B, \mathbf{p}'_B, \mathbf{r}_B) D_\nu(\mathbf{p}, \mathbf{r})^{i\zeta} e(\theta k D_\nu(\mathbf{p}, \mathbf{r})) \mathcal{I}_\nu(x, y, \mathbf{r}), \tag{5.3}$$

where $\mathcal{I}_\nu$ is the indicator of the conditions

$$z' = z + kD_\nu(\mathbf{p}, \mathbf{r}), \qquad p'_{g,i} = p_{g,i} \ (i \in S_{\nu,g}), \qquad r_{g,i} \neq p_{g,j}, p'_{g,j} \quad \text{for all applicable } g, i, j.$$

Each sum is over the full ordered free-array space. Average the source and target slot permutations independently and uniformly, and sum over ν. The resulting row operator is $\mathcal{A}_k^{\zeta,\theta}$. A row kernel acts by $(Af)(x) = \sum_y K(x, y) f(y)$. For $z \in \mathbb{Z}$, let $u_z = \mathbf{1}_{\{\text{position}=z\}}$ in the physical Hilbert space. Equivalently, in ordinary Euclidean coordinates, u_z has value $\sqrt{\tau}$ on that fibre.

5.2. **The residual ideal estimate.** Split $\mathcal{B} = \mathcal{B}_1 \sqcup \mathcal{B}_2$ into two blocks whose cardinalities differ by at most one. Once $m \geq 1$ and $J \geq m$, let I_1 and I_2 be the first m slots in the groups of the respective blocks. The pattern index set is

$$\{0\} \ \sqcup \ \{(A, B) \colon \varnothing \neq A \subseteq I_1, \ \varnothing \neq B \subseteq I_2\}.$$

The raw pattern 0 has $C_{0,g} = \varnothing$ and coefficient 1. The pattern (A, B) replaces $A \cup B$. Thus $|C_{\nu,g}| \leq m$. Its coefficient will be

$$c_{A,B}(\mathbf{p}, \mathbf{p}', \mathbf{r}) = -(-1)^{|A|+|B|} \mathcal{K}_A(P_A(\mathbf{r}), P_A(\mathbf{p}')) \mathcal{K}_B(P_B(\mathbf{r}), P_B(\mathbf{p})), \tag{5.4}$$

where P_A and P_B mean the indicated products.

Choose positive exponents q_*, h_*, x_*, and put

$$Q_0 = \lceil L^{q_*} \rceil, \qquad h = L^{-h_*}, \qquad \xi = L^{-x_*}.$$

For every nonempty selected coordinate set A, define

$$j(v) = \left\lceil \frac{\log v}{h} \right\rceil, \qquad \nu_{A,j} = \Pr(j(P_A) = j).$$

The cells are $((j-1)h, jh]$ in logarithmic coordinates. Let $\mathcal{C}(Q_0)$ be the primitive Dirichlet characters with conductor at most Q_0, including the character of conductor one. Set

$$\mathcal{K}_A(D, X) = \begin{cases} \dfrac{1}{\nu_{A,j(X)}} \displaystyle\sum_{\chi \in \mathcal{C}(Q_0)} \overline{\chi(D)} \chi(X), & j(D) = j(X), \quad \nu_{A,j(X)} \geq \xi, \\ 0, & \text{otherwise.} \end{cases} \tag{5.5}$$

This definition uses the probabilities of the full ordered product spaces.

Lemma 5.1 (Cell comparison)**.** *Fix m, $A_0 > 0$, and $z_0 \geq 0$. The exponents q_*, h_*, x_* can be chosen before J so that the following holds for all sufficiently large L. For nonempty $A \subseteq I_1$, $B \subseteq I_2$, let $\mathbf{X}_A, \mathbf{X}_B, \mathbf{D}_A, \mathbf{D}_B$ be independent tuples with their original product laws. Write $X_A = P_A(\mathbf{X}_A)$, $X_B = P_B(\mathbf{X}_B)$, $D_A = P_A(\mathbf{D}_A)$, and $D_B = P_B(\mathbf{D}_B)$. For arbitrary tests f and g on the full selected tuple spaces of B and A, every $\alpha \in \mathbb{R}$, and $|\zeta| \leq L^{z_0}$,*

$$\Big|\mathbb{E}\overline{f(\mathbf{X}_B)}g(\mathbf{X}_A)\{(X_AX_B)^{i\zeta}e(\alpha X_AX_B) - \mathcal{K}_A(D_A, X_A)\mathcal{K}_B(D_B, X_B)(D_AD_B)^{i\zeta}e(\alpha D_AD_B)\}\Big| \leq L^{-A_0}\|f\|_2\|g\|_2. \tag{5.6}$$

The tests may depend on every coordinate of the indicated tuples.

Proof. Write $v_- = 1/2$, $C_0 = 7/100$, and

$$\kappa = C_0\{\log(m!) + m\log\max(1, v_-^{-1})\}.$$

Unique factorization gives

$$\Pr(P_A = n) \leq \frac{L^\kappa}{n}. \tag{5.7}$$

Indeed, within a group the number of orderings of a prescribed multiset is at most $m!$, and the reciprocal product contributes exactly $1/n$. The normalizing factors cost at most $\max(1, v_-^{-1})^m$ per group. Disjoint prime groups identify the group of each factor. Consequently, for a test f supported on $U \leq P_A < 2U$,

$$\sum_n |\mathbb{E}(f\mathbf{1}_{P_A=n})|^2 \leq \frac{L^\kappa}{U}\|f\|_2^2. \tag{5.8}$$

This is Cauchy–Schwarz separately on each product fibre.

We record the character estimate used in the cells. If $\chi \neq \chi'$ are primitive of conductor at most Q_0, their quotient on a common modulus is nonprincipal. Freeze all selected labels except a prime p, and call the product of the frozen labels R_0. The condition $j(pR_0) = j$ cuts out the prime interval

$$\Big(\max\{a_g, e^{(j-1)h}/R_0\}, \min\{b_g, e^{jh}/R_0\}\Big].$$

Partial summation in the reciprocal-prime form of Lemma 3.5 shows that, for every fixed $D_0 > 0$, the unnormalized off-diagonal cell Gram entry is at most $\rho_0 = L^{-D_0}$, uniformly in the cell and in the frozen labels. All the relevant moduli are bounded by a fixed power of L. For large L, the group primes are coprime to these moduli, so each diagonal Gram entry is exactly $\nu_{A,j}$.

There are at most Q_0^2 characters in $\mathcal{C}(Q_0)$. On a cell of mass at least ξ, the normalized Gram matrix therefore has norm at most $1 + Q_0^2\rho_0/\xi$. Choose D_0 so that $Q_0^2\rho_0 \leq \xi$. The integral operator with kernel $\mathcal{K}_A$ has norm at most 2. Its transpose has the same bound. Distinct cells are orthogonal. Directly from the definition,

$$|\mathcal{K}_A(D, X)| \leq Q_0^2\xi^{-1}, \qquad \mathbb{E}_D|\mathcal{K}_A(D, X)| \leq Q_0^2, \tag{5.9}$$

$$|\mathbb{E}_D\mathcal{K}_A(D, X)\chi(D) - \chi(X)| \leq Q_0^2\xi^{-1}\rho_0 \tag{5.10}$$

on a retained cell. The last estimate also holds for a character modulo $r \leq Q_0$, by passing to its primitive inducer.

The selected products have at most $O_m(L^{d+1})$ dyadic ranges. Fix ranges $X_A \asymp U$, $X_B \asymp V$, and a parameter $S_0 = L^{s_*}$. Dirichlet approximation gives

$$\alpha = a_0/r + \beta, \qquad (a_0, r) = 1, \qquad r \leq UV/S_0, \qquad |\beta| \leq r^{-2}, \qquad |\beta|UV \leq 2S_0/r.$$

For $r > Q_0$, the elementary bilinear geometric-sum estimate, together with (5.8), bounds the contribution by

$$C_m L^{\kappa} \left(\frac{1}{V} + \frac{1}{r} + \frac{1+\log(2r)}{U} + \frac{r(1+\log(2r))}{UV} \right)^{1/2} \|f\|_2 \, \|g\|_2 \, .$$

To obtain this estimate, apply Cauchy–Schwarz in one integer variable and sum its geometric progression. Split the difference variable into intervals of length at most $r/2$. The frequencies in each interval are $1/(2r)$-separated, and their contribution is $O(U + r\log(2r))$. For the term containing the two cell kernels, move each kernel onto its test first. Their operator norms are at most 2, and matched cells enlarge a dyadic support by at most a fixed factor when $h \le \log 2$. Formula (5.8) then gives the same bound.

For $r \le Q_0$, expand $e(a_0 n/r)$ on the units modulo r in multiplicative characters. Parseval bounds the sum of the absolute values of its coefficients by $\sqrt{\phi(r)} \le \sqrt{Q_0}$. Discarded cells have total probability at most

$$C_m L^{d+2+h_*} \xi.$$

Cauchy–Schwarz bounds their contribution by the square root of this quantity times $\|f\|_2 \, \|g\|_2$. On retained cells the two logarithmic products change by at most $2h$. Replacing the phase at the free product by its value at the test product costs at most

$$2h\big(L^{z_0} + 16\pi S_0\big) Q_0^4.$$

This follows by differentiating $\exp(i\zeta u) e(\beta e^u)$ and using (5.9). Formula (5.10) then gives a further error at most

$$3\sqrt{Q_0}\, Q_0^2 \xi^{-1} \rho_0.$$

These pointwise errors pair with the tests at cost $\|f\|_2 \, \|g\|_2$.

For clarity, all choices can be made successively. Put $D = A_0 + 2(d+1) + 1$, and choose

$$q_* = 2(D + \kappa + 2), \qquad s_* = d + 2 + q_*,$$

$$h_* = z_0 + s_* + 4(q_* + 1) + D + 4, \qquad x_* = d + 2 + h_* + 2D + 4.$$

Choose $D_0 > x_* + 3(q_* + 1) + D + 4$. The minor-arc bound, the discarded-cell bound, the phase error and the reproduction error are $O_m(L^{-D})$ on each pair of dyads. The margins absorb their fixed constants. Summing the $O_m(L^{2(d+1)})$ pairs proves (5.6). □

Theorem 5.2 (Residual ideal estimate)**.** *For every $D > 0$, there are $m \ge 1$ and positive q_*, h_*, x_*, C_1, all independent of J, with the following property. For every fixed $J \ge m$, the coefficients (5.4) satisfy, for sufficiently large L,*

$$\sup_{\Theta \in \mathbb{R}} \|\mathcal{T}_\zeta(\Theta)\| \le L^{-D} \qquad (|\zeta| \le L^2), \qquad \sum_{\nu} \|c_\nu\|_\infty \le L^{C_1}. \tag{5.11}$$

One may take

$$C_1 = \frac{7m}{100} \log 2 + 4(q_* + 1) + 2x_* + 1.$$

The partition and coefficient family are fixed before ζ and Θ.

Proof. Let E_i average coordinate i, and let E_{tail} average slot M in every group. For a symmetric vector F, write $F' = E_{\mathrm{tail}} F$, and define H' similarly. Put

$$\Pi_j = \prod_{i \in I_j} (1 - E_i), \qquad j = 1, 2.$$

In a group, consider the m probe coordinates together with its tail coordinate. The $m+1$ projections

$$E_i \prod_{j \ne i} (1 - E_j)$$

are pairwise orthogonal. Symmetry makes the squared norms of their images of F equal. Applying this observation successively in the disjoint groups gives

$$\|\Pi_1 F'\|_2 \le (m+1)^{-|\mathcal{B}_1|/2}\,\|F\|_2\,, \qquad \|\Pi_2 H'\|_2 \le (m+1)^{-|\mathcal{B}_2|/2}\,\|H\|_2\,. \tag{5.12}$$

The remaining tail averages are commuting contractions.

Let $\mathcal{R}(F',H')$ be the raw pairing, including its multiplicative and additive phases. Its norm is at most one. In the pairing for (A,B), the source coordinates in A occur only in F', and the target coordinates in B occur only in H'. Integrating them replaces the tests by $E_A F'$ and $E_B H'$. Fix the shared first-J labels outside $A\cup B$, whose product is C. The remaining source, target and free tuples have precisely the product laws in Lemma 5.1. Apply that lemma with $\alpha=\Theta C$, and multiply by $C^{i\zeta}$. Cauchy–Schwarz in the fixed shared labels gives an error at most $L^{-A_0}\,\|F\|_2\,\|H\|_2$.

This identification uses the full arrays. A label alphabet can equivalently be embedded in the ambient coordinate alphabet by extending every test by zero. The weighted inner product is preserved by this extension, as is integration in each coordinate. Inverse slot permutations preserve the product measure. Thus averaging the two endpoint permutations gives exactly the orthogonal projection S at each end of the stated ideal row operator.

The commuting-projection identities

$$\sum_{\varnothing\neq A\subseteq I_1}(-1)^{|A|}E_A=\Pi_1-1, \qquad \sum_{\varnothing\neq B\subseteq I_2}(-1)^{|B|}E_B=\Pi_2-1$$

now give

$$\mathcal{R}(F',H')-\sum_{A,B\neq\varnothing}(-1)^{|A|+|B|}\mathcal{R}(E_AF',E_BH')=\mathcal{R}(\Pi_1F',H')+\mathcal{R}((1-\Pi_1)F',\Pi_2H').$$

Since $1-\Pi_1$ is an orthogonal projection, the norm of the constructed residual operator is at most

$$(m+1)^{-|\mathcal{B}_1|/2}+(m+1)^{-|\mathcal{B}_2|/2}+2^{ms_B}L^{-A_0}.$$

For sufficiently large L, each block has at least $(7/600)\log L$ groups. Choose m so that $(7/1200)\log(m+1)>D+2$, and then choose $A_0>D+(7m/100)\log 2+2$. Lemma 5.1 supplies its remaining parameters with $z_0=2$. The displayed norm is at most L^{-D} after increasing the threshold.

Finally, (5.9) and the number of selected subsets give

$$\sum_{\nu}\|c_\nu\|_\infty\le 1+2^{ms_B}Q_0^4\xi^{-2}\le L^{(7m/100)\log 2+4(q_*+1)+2x_*+1}.$$

Every choice made here precedes J. □

5.3. **Finite memory estimates.** We give the quantitative estimates used to pass from the ideal operator to physical moments. From now on set

$$R=\lceil L^{13/25}\rceil, \qquad n=2R,$$

and define

$$d_j=\mathbf{1}_{j<n}+\mathbf{1}_{0<j}, \qquad \eta_j=1-q^{d_j/2} \quad (0\le j\le n).$$

For a big prime p, its absent-prime term and tilted mass are

$$b_p=1-\frac{1}{p}\sum_{j=0}^{n}\eta_j, \qquad m_p=\frac{1}{pb_p}, \qquad \nu_g(p)=\frac{m_p}{V_g}, \qquad \beta=\prod_{g\in\mathcal{B}}\prod_{p\in\mathcal{P}_g}b_p. \tag{5.13}$$

For $2(n+1)\le P_*$, these quantities satisfy

$$\frac12\le b_p\le 1, \quad 1\le b_p^{-1}\le 1+\frac{2(n+1)}{P_*}, \quad 1\le\sum_{p\in\mathcal{P}_g}\nu_g(p)\le 2, \quad \nu_g(p)\le\frac{4}{P_*}. \tag{5.14}$$

The finite memory expansion follows from two elementary residue identities. With $H_j(r) = \mathbf{1}_{r+h_j=0 \bmod p}$,

$$\frac{1}{p}\sum_{r \bmod p}\prod_{j=0}^{n}(1-\eta_j H_j(r)) = b_p + \frac{1}{p}\sum_{i<j}\eta_i\eta_j\mathbf{1}_{h_i=h_j \bmod p}\prod_{i<t<j}(1-\eta_t\mathbf{1}_{h_i=h_t \bmod p}). \tag{5.15}$$

If A is a nonempty set of active visits and $a_0 \in A$, then

$$\frac{1}{p}\sum_{r \bmod p}\prod_{j\in A}H_j(r)\prod_{j\notin A}(1-\eta_j H_j(r)) = \frac{\mathbf{1}_{h_j=h_{a_0} \bmod p\ (j\in A)}}{p}\prod_{j\notin A}(1-\eta_j\mathbf{1}_{h_j=h_{a_0} \bmod p}). \tag{5.16}$$

To prove (5.15), apply

$$\prod_j(1-x_j) = 1 - \sum_j x_j + \sum_{i<j}x_i x_j \prod_{i<t<j}(1-x_t)$$

with $x_j = \eta_j H_j(r)$, and then average r. Formula (5.16) follows by fixing $r = -h_{a_0}$.

Process these identities in the order of the visits. A pending particle is a pair (p, r), with its group determined by p. It records an open lifespan and its current residue. Translation adds the complete edge displacement to its residue. The memory is a multiset. A store and a retrieval at an edge of big dilation D_B carry, respectively, the factors

$$\mathbf{1}_{p\nmid D_B}, \qquad V_g^{-1}\mathbf{1}_{p\nmid D_B}.$$

The fresh law is $\nu_g(p)$. Fix the memory damping $\varrho = 1/2$. At visit j, the one-particle ghost coefficients are

$$A_j(p,r) = \frac{1-\eta_j\mathbf{1}_{r=0}}{\varrho^{2\mathbf{1}_{r=0}}}, \qquad D_j(p,r) = -\frac{\eta_j\mathbf{1}_{r=0}}{\varrho}, \qquad C_j(p,r) = -\frac{\eta_j V_{g(p)}\mathbf{1}_{r=0}}{\varrho}.$$

A common cutoff B_0 is imposed at each intermediate memory state, including the state after all stores and before retrieval. At a visit, a pending multiplicity m_u has a retention choice $0 \le r_u \le m_u$. A fresh batch is an ordered tuple $(u_1, \dots, u_b)$. Its contribution to a ghost transition is

$$\left\{\prod_u \binom{m_u}{r_u} A_j(u)^{r_u} D_j(u)^{m_u-r_u}\right\}\frac{1}{b!}\prod_{\ell=1}^{b} C_j(u_\ell)\nu_{g(u_\ell)}(p(u_\ell)).$$

The target multiset is $\sum_u r_u[u] + \sum_{\ell=1}^{b}[u_\ell]$. The order of the fresh tuple specifies its birth addresses. The factors ϱ in the adjacent store blocks cancel the displayed ghost denominators, recovering (5.15) and (5.16).

5.3.1. *The weighted store and retrieval kernels.* We give the coordinate calculation for the memory edge. Throughout this calculation the tilted laws $\nu_g(p) = m_p/V_g$ are those in (5.13). For a big group g, put

$$\Omega_g = \{(p,r)\colon p \in \mathcal{P}_g,\ r \in \mathbb{Z}/p\mathbb{Z}\}, \qquad e_g(p) = (p,0),$$
$$\nu_g(p,r) = \nu_g(p), \qquad N_g = \sum_{p\in\mathcal{P}_g}\nu_g(p).$$

Each residue of a fixed prime has weight $\nu_g(p)$. The factor $1/p$ has already entered m_p. A memory s is a multiset on $\Omega = \bigsqcup_{g\in\mathcal{B}}\Omega_g$. Write s_g for its restriction to Ω_g, and set

$$|s| = \sum_{u\in\Omega}s(u), \qquad Z_g(s) = \sum_{p\in\mathcal{P}_g}s_g(e_g(p)), \qquad Z(s) = \sum_{g\in\mathcal{B}}Z_g(s),$$

$$\pi(s) = \prod_{g\in\mathcal{B}}\prod_{u\in\Omega_g}\frac{\nu_g(u)^{s_g(u)}}{s_g(u)!}. \tag{5.17}$$

For an active big-label array $\mathbf{a}$, the mass of the full mixed state $x = (x_S, \mathbf{a}, s)$ is

$$m_X(x) = \tau_S \left(\prod_{g \in \mathcal{B}} \prod_{i=1}^{M} \nu_g(a_{g,i}) \right) \pi(s). \tag{5.18}$$

Here x_S is the small physical vertex, whose position and labels are part of the state. The cutoff is $|s| \le B_0$. To calculate the local kernels we first allow all active arrays, with their product mass. The original distinct-label condition is imposed by restriction to the corresponding states.

Kernels act by ordinary finite sums,

$$(Kf)(x) = \sum_y K(x,y) f(y), \qquad \langle f, h \rangle_{m_X} = \sum_x m_X(x) \overline{f(x)} h(x).$$

Consequently the weighted adjoint has kernel

$$K^*(x,y) = \frac{m_X(y)}{m_X(x)} \overline{K(y,x)}. \tag{5.19}$$

Fix one group, suppress its subscript, and write π_g for its factor in (5.17). For one active slot, a parameter $\alpha > 0$, and a real coefficient $b(p) \ge 0$, define

$$S_{\alpha,b}((a,s),(a',t)) = \nu(a') \alpha^{Z_g(s)} \left(\mathbf{1}_{t=s} + b(a) \mathbf{1}_{t=s+[e(a)]} \right). \tag{5.20}$$

Both memories in this formula belong to the cutoff space. Thus a store that would exceed the cutoff contributes zero. On a complete array, this kernel also contains the indicator that every other active slot is fixed. The term $t = s$ drops the old active label. The term $t = s + [e(a)]$ stores that label at residue zero. The new active label a' has weight $\nu(a')$ in either case.

The elementary factorial identity

$$\pi_g(s + [u])(s(u) + 1) = \nu(u) \pi_g(s) \tag{5.21}$$

determines the reverse kernel. With the one-slot mass $\nu(a)\pi_g(s)$, formula (5.19) gives

$$\begin{aligned} S^*_{\alpha,b}((a,s),(a',t)) &= \nu(a') \alpha^{Z_g(s)} \mathbf{1}_{t=s} \\ &\quad + b(a') s(e(a')) \alpha^{Z_g(t)} \mathbf{1}_{s=t+[e(a')]}. \end{aligned} \tag{5.22}$$

Indeed, the drop terms on the two sides of weighted adjointness are both $\nu(a)\nu(a')\pi_g(s)\alpha^{Z_g(s)}$. For a store term, (5.21) replaces $\nu(a')\pi_g(t)$ by $\pi_g(s)s(e(a'))$, giving the second term of (5.22). The masses of all the other active slots cancel. This proves the formula on the full array space as well.

Thus retrieval either samples a fresh active label with weight $\nu(a')$, or removes one pending zero-residue copy of the target label with weight $b(a')s(e(a'))$. The latter multiplicity is evaluated in the memory immediately before that removal. This description includes repeated numerical labels and arbitrary pending multiplicities.

5.3.2. *A block with one damping factor.* Let $i_1, \dots, i_t$ be the unshared slots of one group in their original increasing order. They are distinct. Let $S_{i,\alpha,b}$ denote the full-array version of (5.20), and let

$$(D_\lambda f)(\mathbf{a}, s) = \lambda^{Z_g(s)} f(\mathbf{a}, s).$$

The source block in an edge is exactly

$$\mathcal{S}_{\varrho,\mathbf{b}} = D_\varrho S_{i_1,1,b_1} \cdots S_{i_t,1,b_t}, \qquad \varrho = \tfrac{1}{2}. \tag{5.23}$$

It applies one factor $\varrho^{Z_g(s)}$, evaluated in the source memory. Its adjoint is

$$\mathcal{S}^*_{\varrho,\mathbf{b}} = S^*_{i_t,1,b_t} \cdots S^*_{i_1,1,b_1} D_\varrho. \tag{5.24}$$

The target slots are therefore processed in decreasing original order, and the damping factor is evaluated in the final pending memory. Both formulas retain the cutoff after each individual store or removal.

Here is the exact distribution of the single damping factor used in the estimates. For $0 < \vartheta < 1$ and an integer $r \geq 0$, the two alternatives in (5.20) give

$$S_{i,\vartheta,b\vartheta^{-r}} D_{\vartheta^r} = D_{\vartheta^{r+1}} S_{i,1,b}. \tag{5.25}$$

For a drop, $Z_g(t) = Z_g(s)$. For a store, $Z_g(t) = Z_g(s) + 1$, and the additional ϑ^r is canceled by ϑ^{-r}. Multiplying (5.25) for the successive slots yields

$$\prod_{j=1}^{t} S_{i_j,\vartheta,b_j\vartheta^{-(t-j)}} = D_{\vartheta^t} \prod_{j=1}^{t} S_{i_j,1,b_j}. \tag{5.26}$$

Products of operators here and below are ordered from left to right in increasing j, as in the row action of the chronological kernel. If $t \leq T_0$ and $\varrho \leq \vartheta^{T_0}$, then $0 < \varrho/\vartheta^t \leq 1$. Hence (5.23) equals the left side of (5.26) preceded by the diagonal contraction D_{ϱ/ϑ^t}.

The one-slot estimates follow directly from (5.20) and (5.22). Use the positive weight $w(s) = v^{|s|}$, with $v = 16$. For $b = 1$, the weighted row sum is at most $N_g(1 + v)w(s)$. The weighted column sum, divided by the target mass and $w(t)$, is at most

$$N_g \vartheta^{Z_g(t)} + \frac{Z_g(t)\vartheta^{Z_g(t)-1}}{v} \leq N_g + \frac{1}{v(1-\vartheta)}.$$

The removal term is zero when $Z_g(t) = 0$. For a positive integer k, its bound follows from

$$k\vartheta^{k-1} \leq \sum_{j=0}^{k-1} \vartheta^j \leq \frac{1}{1-\vartheta}.$$

If $0 \leq b \leq A$ and $A \geq 1$, both estimates acquire a factor at most A, by termwise comparison with the kernel for $b = 1$. These estimates are independent of the number of other active slots and of the cutoff.

Take $T_0 = m + 1$, $\vartheta = 1 - 1/(2T_0)$, and $\Lambda = \max_p b_p^{-1}$, with the notation of the complete-edge estimate below. Then $1 \leq N_g \leq \Lambda$ and $\vartheta^{T_0} \geq 1/2 = \varrho$. Indeed, $\nu_g(p) = b_p^{-1}\mu_g(p)$, and Bernoulli's inequality gives $(1 - 1/(2T_0))^{T_0} \geq 1 - T_0/(2T_0) = 1/2$. If $0 \leq b_j \leq 2$, each compensating coefficient in (5.26) is at most $2\vartheta^{-T_0}$. The source block has weighted row and column bounds R_0^t and C_0^t, where

$$R_0 = 2\vartheta^{-T_0}\Lambda(1 + 16), \qquad C_0 = 2\vartheta^{-T_0}\left(\Lambda + \frac{1}{16(1-\vartheta)}\right). \tag{5.27}$$

Its weighted adjoint has row bound C_0^t and column bound R_0^t.

5.3.3. *Translation and the auxiliary-label normalization.* Fix an integer displacement δ. Translation acts on each particle by

$$T_\delta(p, r) = (p, r + \delta \bmod p), \qquad (U_\delta f)(\mathbf{a}, s) = f(\mathbf{a}, T_\delta s).$$

This map preserves $|s|$ and $\pi(s)$, since the particle weight depends only on its prime. It therefore preserves the weighted norm and the weight $16^{|s|}$.

For one group define the normalized memory edge by

$$\mathcal{E}_g = N_g^{-t} \mathcal{S}_{\varrho,\mathbf{b}_s} U_\delta \mathcal{S}^*_{\varrho,\mathbf{b}_r}. \tag{5.28}$$

Since $U_\delta^* = U_{-\delta}$, this definition gives the exact identity

$$\mathcal{E}_g(\mathbf{b}_s, \mathbf{b}_r, \delta)^* = \mathcal{E}_g(\mathbf{b}_r, \mathbf{b}_s, -\delta). \tag{5.29}$$

The factor N_g^{-t} has an exact finite-sum origin. In the source block each selected slot is replaced by an auxiliary active label of weight ν_g. These auxiliary labels range over the full product alphabet. The physical distinctness restrictions are applied at the original source and final target of the edge. The retrieval kernel (5.22) is independent of its input active label. Thus the auxiliary label at each selected slot is summed freely, giving

$$\sum_{d_1,\dots,d_t\in\mathcal{P}_g}\prod_{j=1}^{t}\nu_g(d_j)=N_g^t.$$

This cancels the prefactor in (5.28). The distinctness of $i_1,\dots,i_t$ ensures that every source store reads its original active label. The memory cutoff depends only on the store mask and the source memory. The final coefficient depends on the final target array and memory. Both therefore remain fixed during the auxiliary-label sums. This proves the cancellation with an arbitrary coefficient on the complete final state.

More explicitly, choose a store mask $\epsilon\in\{0,1\}^t$ and put

$$s_\epsilon=s+\sum_{j=1}^{t}\epsilon_j[e(a_{i_j})].$$

This choice contributes $\prod_{j:\epsilon_j=1}b_{s,j}(a_{i_j})$, subject to $|s_\epsilon|\le B_0$. Start the reverse retrieval at $u_t=T_\delta s_\epsilon$. For $j=t,t-1,\dots,1$, fix the target label a'_{i_j} and choose a fresh flag f_j. Set

(5.30)

f_j	u_{j-1}	q_j
1	u_j	$\nu_g(a'_{i_j})$
0	$u_j-[e(a'_{i_j})]$	$b_{r,j}(a'_{i_j})u_j(e(a'_{i_j}))$.

Removal of an unavailable copy gives zero weight. The terminal memory is u_0. The edge kernel is exactly

$$\begin{aligned}\mathcal{E}_g((\mathbf{a},s),(\mathbf{a}',s'))&=\mathbf{1}_{a_i=a'_i\ (i\notin\{i_1,\dots,i_t\})}\varrho^{Z_g(s)+Z_g(s')}\\&\quad\times\sum_{\epsilon\in\{0,1\}^t}\mathbf{1}_{|s_\epsilon|\le B_0}\left(\prod_{j:\epsilon_j=1}b_{s,j}(a_{i_j})\right)\sum_{\mathbf{f}\in\{0,1\}^t}\mathbf{1}_{u_0=s'}\prod_{j=1}^{t}q_j.\end{aligned}\tag{5.31}$$

Expanding the source kernels, using (5.22) in reverse order, and summing the auxiliary active labels proves (5.31). All sums are finite, so this calculation also applies after multiplication by arbitrary complex endpoint coefficients.

For several big groups use the product of these kernels with the single total cutoff $\sum_g|s_g|\le B_0$. Equivalently, insert the corresponding coordinate projection before the source block, at its output, and after the retrieval block. During source storage the total memory size is nondecreasing, so its largest value is the fully stored value. Translation preserves this value and every retrieval decreases it or leaves it fixed. Thus the cutoff in the multigroup version of (5.31) is exactly

$$\left|s+\sum_{g\in\mathcal{B}}\sum_{j=1}^{t_g}\epsilon_{g,j}[e_g(a_{g,i_j})]\right|\le B_0.\tag{5.32}$$

This also enforces every intermediate source-store cutoff, even when a later retrieval would lower the memory size.

The weighted row and column estimates compose through translation. Since $N_g^{-t_g}\le 1$, (5.27) gives a common bound $(R_0C_0)^{t_g}$ for both sums in group g. Restricting the positive

memory kernels to the total-cutoff states preserves these bounds. Their product over the big groups is at most

$$\left[\max\left\{1,(2\vartheta^{-T_0})^2\Lambda(1+16)\left(\Lambda+\frac{1}{16(1-\vartheta)}\right)\right\}\right]^{T_0 s_B}=C_{\mathrm{mem}}.$$

This proves the stated memory constant with one source damping factor and one final retrieval damping factor per edge.

5.3.4. *The original mixed edge.* Fix an original choice consisting of a pattern ν, its full ordered free tuple $\mathbf{r}$, and the independent source and target slot permutations σ,τ. Write $\mathbf{a}=\mathbf{p}^{\sigma}$, $\mathbf{a}'=\mathbf{p}'^{\tau}$, $D=D_\nu(\mathbf{a},\mathbf{r})$, and let D_B be its big part. For the forward edge, $\delta=kD$. The exact store and retrieval coefficients in a big group are

$$b_{s,g,j}(p)=\mathbf{1}_{p\nmid D_B},\qquad b_{r,g,j}(p)=V_g^{-1}\mathbf{1}_{p\nmid D_B}. \tag{5.33}$$

They satisfy $0\le b_s\le 1$ and $0\le b_r\le 2$. The dilation uses only shared source slots and free labels. Consequently D, D_B, and the translation are fixed while the unshared slots are processed.

Because the prime groups are disjoint and $p\nmid k$, $p\nmid D_B$ implies $\delta\not\equiv 0 \pmod p$. A newly stored particle $(p,0)$ therefore becomes (p,δ) with nonzero residue. This proves directly that it contributes zero to every zero-residue retrieval in the same edge. On an endpoint satisfying the original distinct-label, shared-label and free-label guards, each unshared source or target prime is coprime to D_B. The coefficients in (5.33) then give a source-store factor 1, a fresh-target weight m_p/V_g, and a retrieval weight V_g^{-1} for each available pending copy.

Let $\Pi=\prod_g \mathrm{Sym}([M])$, and evaluate $\mathcal{I}_\nu$ using the displayed label and integer conditions in (5.3). For this fixed forward choice, the scalar multiplying the memory kernel is

$$\begin{aligned}\kappa_{\nu,\sigma,\tau,\mathbf{r}}(x,y)=\frac{d_{q,S}(x_S)d_{q,S}(y_S)}{|\Pi|^2}\left(\prod_{g\in\mathcal{S}}V_g^{-t_{\nu,g}}\right)\left(\prod_g\prod_{i\in C_{\nu,g}}\mu_g(r_{g,i})\right)\\ \times c_\nu(\mathbf{a}_B,\mathbf{a}'_B,\mathbf{r}_B)D^{i\zeta}e(\theta kD)\,\mathcal{I}_\nu(x^\sigma,y^\tau,\mathbf{r}).\end{aligned} \tag{5.34}$$

Here $d_{q,S}$ denotes the damping from the small groups, and the guard includes the exact integer equation $\mathrm{position}(y)=\mathrm{position}(x)+kD$. The complete forward mixed edge is the finite sum of (5.34) times the multigroup kernel (5.31), over precisely these original choices. Shared labels are fixed by the kernel. Restricting its endpoint arrays imposes the original distinct-label conditions. The full free tuple and both independent permutation averages remain in this formula.

Every changed big target slot already contributes its inverse group mass through its fresh weight $\nu_g=m_p/V_g$ or its retrieval coefficient V_g^{-1}. This accounts for the big target normalization removed in (5.34). Similarly, the initial active product in (5.18) supplies the tilted initial factors. Expanding the finite sums establishes this identity before taking absolute values of any signed pattern coefficient.

For a reversed edge the scalar is the complex conjugate of the forward scalar with the endpoints interchanged. The memory translation becomes $T_{-\delta}$, the source and target permutations interchange, and the store and retrieval coefficient families interchange. Formula (5.19) then proves the exact weighted adjoint identity on the mixed state space. In particular, the reverse row estimate is the required mass-weighted column estimate for the forward edge.

Finally, the placement of the two damping factors agrees with the ghost coefficients. At a zero-residue visit,

$$\varrho^2A_j(p,0)=1-\eta_j,\qquad \varrho D_j(p,0)=-\eta_j,\qquad \varrho C_j(p,0)\nu_g(p)=-\eta_j m_p.$$

A retained pending copy receives the factors from the adjacent incoming and outgoing edges. A destroyed copy receives the incoming factor, and a newborn copy receives the outgoing factor. Off zero, every such damping factor is one. These identities recover the corresponding lifespan terms with the initial and terminal pending memories empty.

For use below, let E denote one complete edge on this weighted memory space, including the small physical labels, both endpoint permutations, the original free-array sum and its coefficients. Let $E(\beta_s,\beta_r)$ display its store and retrieval factors. Write

$$C = E(0,0), \qquad P = E(\beta_s,\beta_r) - E(\beta_s,0), \qquad S_1 = E(\beta_s,0) - E(0,0).$$

These are classes of traces on the same space, and $E = C + P + S_1$. Define

$$\Lambda = \max_{p \in \bigcup_{g\in\mathcal{B}} \mathcal{P}_g} b_p^{-1}, \quad K = \sum_{\nu} \|c_\nu\|_\infty, \quad t = (m+1)s_B,$$

$$C_{\mathrm{cond}} = \prod_{g\in\mathcal{S}} \frac{1}{(1-q)V_g}, \qquad C_{\mathrm{row}} = \prod_{g\in\mathcal{S}} \frac{1}{(1-\sqrt{q})V_g}.$$

Put $T_0 = m+1$, $\vartheta = 1 - 1/(2T_0)$, and

$$\mathfrak{C}(\Lambda,\vartheta) = \left[\max\left\{1,\ (2\vartheta^{-T_0})^2\Lambda(1+16)\left(\Lambda + \frac{1}{16(1-\vartheta)}\right)\right\}\right]^{T_0}, \qquad C_{\mathrm{mem}} = \mathfrak{C}(\Lambda,\vartheta)^{s_B}.$$

Lemma 5.3 (Complete-edge estimate)**.** *Suppose $\|\mathcal{T}_\zeta(\Theta)\| \le \epsilon$ for every $\Theta \in \mathbb{R}$, and $p \nmid k$ for every big prime. For the actual two-scale groups, sufficiently large L gives*

$$\|E\| \le C_{\mathrm{cond}}\left\{\epsilon + K(\Lambda^t - 1) + K\Lambda^t b_{\mathrm{col}}\right\} + 4KC_{\mathrm{row}}C_{\mathrm{mem}}\sqrt{\frac{B_0}{(J+1)^{s_S}}}, \tag{5.35}$$

$$b_{\mathrm{col}} = \frac{2m(J+m+2)s_B}{P_*}. \tag{5.36}$$

The estimate holds for both orientations and for every finite position window.

Proof. The weighted completion map.

We describe the clean edge in the notation of Lemma 5.3. Here $M = J+1$, the fresh law is $\nu_g(p) = \mu_g(p)b_p^{-1}$, and the memory cutoff is B_0. Let

$$\mathcal{A}_B = \prod_{g\in\mathcal{B}} \mathcal{P}_g^M, \qquad w_B(\mathbf{a}) = \prod_{g\in\mathcal{B}}\prod_{i=1}^{M} \nu_g(a_{g,i}),$$

and let $\mathcal{A}_B^{\mathrm{inj}}$ be the subset on which the entries in each group are distinct. A memory $\mathfrak{m}$ is a multiset of pairs (p,r), with p in a big group and $r \in \mathbb{Z}/p\mathbb{Z}$, whose cardinality is at most B_0. Write

$$\pi(\mathfrak{m}) = \prod_{g\in\mathcal{B}}\prod_{p\in\mathcal{P}_g}\prod_{r \bmod p} \frac{\nu_g(p)^{\mathfrak{m}(p,r)}}{\mathfrak{m}(p,r)!}, \qquad Z(\mathfrak{m}) = \sum_{g\in\mathcal{B}}\sum_{p\in\mathcal{P}_g} \mathfrak{m}(p,0).$$

The small physical space $\mathcal{V}_S$ consists of $u = (z, \mathbf{p}_S)$, where $z \in \mathcal{W}$, all the labels divide z, and the M labels in each small group are distinct. Put

$$\tau_S = \prod_{g\in\mathcal{S}} V_g^{-M}, \qquad \alpha_S = \prod_{g\in\mathcal{S}} V_g^{-1}, \qquad d_S(u) = q^{(\omega_S(z) - Ms_S)/2}.$$

Thus the joint state space and its mass are

$$\mathcal{X} = \mathcal{V}_S \times \mathcal{A}_B^{\mathrm{inj}} \times \mathcal{M}_{B_0}, \qquad m_X(u,\mathbf{a},\mathfrak{m}) = \tau_S w_B(\mathbf{a})\pi(\mathfrak{m}).$$

All inner products below are conjugate-linear in the first variable. In particular,

$$\langle f, g\rangle_X = \sum_{u,\mathbf{a},\mathfrak{m}} \tau_S w_B(\mathbf{a})\pi(\mathfrak{m})\overline{f(u,\mathbf{a},\mathfrak{m})}g(u,\mathbf{a},\mathfrak{m}).$$

Record the first J small labels by

$$\mathbf{d}(u) = (p_{g,i})_{g\in\mathcal{S},\ 1\le i\le J}, \qquad \Xi = \prod_{g\in\mathcal{S}} \mathcal{P}_g^J, \qquad d(\mathbf{d}) = \prod_{g\in\mathcal{S}} \prod_{i=1}^{J} d_{g,i}.$$

The ambient prefix space Ξ contains every ordered tuple.

We choose an exact position circle before defining the target space. Let Δ be the finite set of all integers $\pm k D_\nu(\mathbf{p},\mathbf{r})$, with ν ranging over the complete pattern family and $\mathbf{p},\mathbf{r}$ over their full ordered label spaces. Set

$$N = 1 + \max_{z,z'\in\mathcal{W},\ \delta\in\Delta\cup\{0\}} |z' - z - \delta|,$$

where the maximum of an empty set is zero. For every relevant edge,

$$z' = z + \delta \pmod N \quad\Longleftrightarrow\quad z' = z + \delta. \tag{5.37}$$

Indeed, the difference is divisible by N and has absolute value less than N. The case $\delta = 0$ also proves that $\iota\colon \mathcal{W} \to \mathbb{Z}/N\mathbb{Z}$, $z \mapsto z \bmod N$, is injective. The set Δ includes the arrays introduced by zero extension and both endpoint permutation averages.

The completion target is

$$\mathcal{T} = \Xi \times \mathbb{Z}/N\mathbb{Z} \times \mathcal{A}_B \times \mathcal{M}_{B_0}, \qquad m_T(\mathbf{d},\bar z,\mathbf{a},\mathfrak{m}) = w_B(\mathbf{a})\pi(\mathfrak{m}).$$

Prefix and position carry counting measure. Define

$$(\mathcal{J}f)(\mathbf{d},\bar z,\mathbf{a},\mathfrak{m}) = \sqrt{\tau_S\alpha_S} \sum_{\substack{u\in\mathcal{V}_S:\ \mathbf{d}(u)=\mathbf{d}\\ \iota(z(u))=\bar z}} d_S(u) f(u,\mathbf{a},\mathfrak{m}) \tag{5.38}$$

when $\mathbf{a} \in \mathcal{A}_B^{\mathrm{inj}}$, and give it value zero for other $\mathbf{a}$. Empty fibres also have value zero. This formula keeps the small-label divisibility and distinctness conditions inside the summation. The weighted adjoint is

$$(\mathcal{J}^*F)(u,\mathbf{a},\mathfrak{m}) = \sqrt{\frac{\alpha_S}{\tau_S}}\, d_S(u) F(\mathbf{d}(u),\iota(z(u)),\mathbf{a},\mathfrak{m}). \tag{5.39}$$

Indeed, direct expansion gives

$$\sum_{t\in\mathcal{T}} m_T(t)\overline{(\mathcal{J}f)(t)}F(t) = \sum_{x\in\mathcal{X}} m_X(x)\overline{f(x)}\sqrt{\frac{\alpha_S}{\tau_S}} d_S(u) F(\operatorname{tag} x),$$

since $m_X = \tau_S m_T$ on the fibre of the displayed tag. In particular, if $U(t,t')$ is a row kernel on $\mathcal{T}$, then

$$(\mathcal{J}^*U\mathcal{J})(x,x') = \alpha_S d_S(u) d_S(u')\, U(\operatorname{tag} x, \operatorname{tag} x'). \tag{5.40}$$

Thus the two completion factors give exactly the original small target normalization and the two small endpoint damping factors.

For a fixed $(\mathbf{d},\bar z)$, the position z is unique if the fibre is nonempty. Put $n_g = \omega_{\mathcal{P}_g}(z) - J$. There are n_g choices for the final label in group g, and, for an admissible prefix, consequently

$$\sum_{u \text{ in the fibre}} d_S(u)^2 = \prod_{g\in\mathcal{S}} n_g q^{n_g - 1}.$$

For $n \ge 1$,

$$n q^{n-1} \le \sum_{r=0}^{n-1} q^r \le \frac{1}{1-q}.$$

If one of the required choices is absent, the fibre contributes zero. It follows that

$$\alpha_S \sum_{u \text{ in the fibre}} d_S(u)^2 \le \prod_{g\in\mathcal{S}} \frac{1}{(1-q)V_g} = C_{\mathrm{cond}}. \tag{5.41}$$

Cauchy–Schwarz in each fibre, followed by summation with the unchanged weights $w_B(\mathbf{a})\pi(\mathfrak{m})$, now gives

$$\|\mathcal{J}f\|_T^2 \le C_{\mathrm{cond}}\|f\|_X^2, \qquad \|\mathcal{J}\|^2 \le C_{\mathrm{cond}}. \tag{5.42}$$

The exact clean kernel on the completion space.

Let Π_B and Π_S be the products of the slot-permutation groups over the big and small groups, and put $\Pi_G = \Pi_S \times \Pi_B$. For an array $\mathbf{a}$, write $\mathbf{a}^\sigma_{g,i} = a_{g,\sigma_g(i)}$. The independent big endpoint averages remain inside the following kernel. Fix $\mathbf{d} \in \Xi$, and write $d = d(\mathbf{d})$. For a pattern ν, put

$$w_\nu(\mathbf{r}) = \prod_{g\in\mathcal{B}} \prod_{i\in C_{\nu,g}} \mu_g(r_{g,i}), \qquad \nu_{U_\nu}(\mathbf{b}) = \prod_{g\in\mathcal{B}} \prod_{i\in U_{\nu,g}} \nu_g(b_{g,i}).$$

Let $G_\nu(\mathbf{a},\mathbf{b},\mathbf{r})$ be the indicator that all shared big coordinates agree and that every free label differs from every source and target label in its group. It imposes the original free-label exclusions and allows free labels to repeat one another. Let T_h translate a position by h mod N and each residue of a pending particle (p,r) by h mod p, with multiplicities preserved. Define the row kernel

$$\begin{aligned} U_{\mathbf{d}}\big((\bar z,\mathfrak{m},\mathbf{a}),(\bar z',\mathfrak{m}',\mathbf{b})\big) = \frac{1}{|\Pi_B|^2} \sum_{\sigma,\tau\in\Pi_B} \sum_{\nu} \sum_{\mathbf{r}} w_\nu(\mathbf{r})\nu_{U_\nu}(\mathbf{b}^\tau) G_\nu(\mathbf{a}^\sigma,\mathbf{b}^\tau,\mathbf{r}) \\ \cdot c_\nu(\mathbf{a}^\sigma,\mathbf{b}^\tau,\mathbf{r})(dD)^{i\zeta} e(\theta k d D)\mathbf{1}_{(\bar z',\mathfrak{m}')=T_{kdD}(\bar z,\mathfrak{m})}. \end{aligned} \tag{5.43}$$

Here $D = D_{\nu,B}(\mathbf{a}^\sigma,\mathbf{r})$ in each summand. The full operator U is the direct sum of these kernels over $\mathbf{d}$. Its mass on each component is $w_B(\mathbf{a})\pi(\mathfrak{m})$, with counting measure in $\bar z$.

Let P_S average the small-label slot permutations on $\mathcal{X}$, and let $(D_\varrho f)(x) = \varrho^{Z(\mathfrak{m})} f(x)$. Both are contractions for the displayed mass. The exact clean identity is

$$C = D_\varrho P_S \mathcal{J}^* U \mathcal{J} P_S D_\varrho. \tag{5.44}$$

To check it, expand the two factors P_S, the two factors in (5.40), and the kernel (5.43). Equality of the prefix indices is the small shared-coordinate condition. The full dilation equals $d(\mathbf{d})D_{\nu,B}$, since small groups have empty replacement sets. Their free-label spaces are singletons and contribute weight one. The big target factor is precisely the ordered fresh law ν_{U_ν}. The two factors from D_ϱ are the original pending-zero damping factors. Formula (5.37) recovers the integer position equation from the torus equation. Finally, $|\Pi_S|^{-2}|\Pi_B|^{-2} = |\Pi_G|^{-2}$, so every original endpoint permutation and its normalization is present. The all-fresh trace leaves the pending cardinality unchanged. Hence its intermediate cutoff is the same B_0 throughout this identity.

An explicit isometric Fourier lift.

Choose a positive integer N_1 divisible by N and by every big prime, for example $N_1 = N\prod_{g\in\mathcal{B}}\prod_{p\in\mathcal{P}_g} p$. Put $A_{N_1} = \mathbb{Z}/N_1\mathbb{Z}$ and $Y = (\mathbb{Z}/N\mathbb{Z}) \times \mathcal{M}_{B_0}$. Then T_a, $a \in A_{N_1}$, is a well-defined action on Y, and

$$T_{a+b} = T_b T_a, \qquad \pi(T_a\mathfrak{m}) = \pi(\mathfrak{m}). \tag{5.45}$$

The second equality follows by permuting the residue classes of each fixed prime. The numerator and the multiplicity factorials are both preserved.

Let $\mathcal{H}_\nu = L^2(\mathcal{A}_B, w_B)$, and regard a function on $Y \times \mathcal{A}_B$ as a function on Y with values in $\mathcal{H}_\nu$. The orbit lift is

$$(\mathcal{L}F)(a,y) = F(T_a y), \qquad a \in A_{N_1}.$$

Give A_{N_1} the probability measure N_1^{-1} and Y the mass $\pi(y) = \pi(\mathfrak{m})$. Direct reindexing gives

$$\frac{1}{N_1} \sum_{a \in A_{N_1}} \sum_{y \in Y} \pi(y) \|F(T_a y)\|_{\mathcal{H}_\nu}^2 = \sum_{y \in Y} \pi(y) \|F(y)\|_{\mathcal{H}_\nu}^2. \tag{5.46}$$

For an $\mathcal{H}_\nu$-valued function F on A_{N_1}, use

$$\widehat{F}(\chi) = \frac{1}{N_1} \sum_{a \in A_{N_1}} \overline{\chi(a)} F(a).$$

For $a, b \in A_{N_1}$, the geometric sum gives

$$\sum_{j=0}^{N_1-1} e\Big(\frac{j(a-b)}{N_1}\Big) = N_1 \mathbf{1}_{a=b}.$$

Expanding squared norms and using this identity gives

$$\sum_{\chi} \|\widehat{F}(\chi)\|_{\mathcal{H}_\nu}^2 = \frac{1}{N_1} \sum_{a} \|F(a)\|_{\mathcal{H}_\nu}^2, \qquad \widehat{F(\cdot + h)}(\chi) = \chi(h) \widehat{F}(\chi). \tag{5.47}$$

These formulas apply to every orbit, with its actual stabilizer and multiplicity, because the calculation sums over all $a \in A_{N_1}$.

Slice the complete signed, independently averaged big-label kernel by its integer dilation D. Denote by B_D the row operator whose kernel is

$$B_D(\mathbf{a}, \mathbf{b}) = \frac{1}{|\Pi_B|^2} \sum_{\sigma, \tau \in \Pi_B} \sum_{\nu} \sum_{\mathbf{r}} w_\nu(\mathbf{r}) \nu_{U_\nu}(\mathbf{b}^\tau) G_\nu(\mathbf{a}^\sigma, \mathbf{b}^\tau, \mathbf{r})\, c_\nu(\mathbf{a}^\sigma, \mathbf{b}^\tau, \mathbf{r})\, D^{i\zeta} \cdot \mathbf{1}_{D_{\nu,B}(\mathbf{a}^\sigma, \mathbf{r}) = D}.$$

Only finitely many positive D occur. Formula (5.43) becomes the operator identity

$$(U_{\mathbf{d}} F)(y) = \sum_{D} d^{i\zeta} e(\theta k d D) B_D F(T_{kdD} y).$$

Applying $\mathcal{L}$, relation (5.45) changes the argument to $a + kdD$. The Fourier block at a character $\chi(a) = e(ja/N_1)$ is therefore exactly

$$\begin{aligned} \mathcal{B}_{\mathbf{d},\chi} &= d^{i\zeta} \sum_{D} e(\theta k d D) \chi(kdD) B_D \\ &= d^{i\zeta}\, \widetilde{\mathcal{T}}_\zeta^{\mathrm{allow}} \Big(kd \Big(\theta + \frac{j}{N_1} \Big) \Big). \end{aligned} \tag{5.48}$$

Here $\widetilde{\mathcal{T}}^{\mathrm{allow}}$ is the literal ideal row operation with fresh law ν_g, the original free exclusions, and both endpoint averages. All pattern terms have already been summed in this block. The scalar $d^{i\zeta}$ has modulus one.

The Fourier components of a physical test are

$$\mathcal{J}_{\mathbf{d},\chi,y} f = \frac{1}{N_1} \sum_{a \in A_{N_1}} \overline{\chi(a)} (\mathcal{J} P_S D_\varrho f)(\mathbf{d}, T_a y, \cdot) \quad \text{in } \mathcal{H}_\nu. \tag{5.49}$$

Equations (5.44)– (5.48) give the exact pairing

$$\langle f, Cg \rangle_X = \sum_{\mathbf{d} \in \Xi} \sum_{\chi} \sum_{y \in Y} \pi(y) \left\langle \mathcal{J}_{\mathbf{d},\chi,y} f, \mathcal{B}_{\mathbf{d},\chi} \mathcal{J}_{\mathbf{d},\chi,y} g \right\rangle_{\mathcal{H}_\nu}. \tag{5.50}$$

Moreover, the exact isometry and Parseval identities give

$$\sum_{\mathbf{d},\chi,y} \pi(y) \|\mathcal{J}_{\mathbf{d},\chi,y} f\|_{\mathcal{H}_\nu}^2 = \|\mathcal{J} P_S D_\varrho f\|_T^2 \le C_{\mathrm{cond}} \|f\|_X^2. \tag{5.51}$$

This is the normalization used at both ends of the pairing.

The tilted law and the original free exclusions.

Put $r_g(p) = b_p^{-1}$, and $r(\mathbf{a}) = \prod_{g,i} r_g(a_{g,i})$. The map

$$(\mathcal{V}F)(\mathbf{a}) = \sqrt{r(\mathbf{a})}F(\mathbf{a})$$

is an isometry from $\mathcal{H}_\nu$ to the harmonic space $\mathcal{H}_B$, and commutes with every slot permutation. On the shared-coordinate support, conjugation by $\mathcal{V}$ changes a pattern's harmonic row kernel by the multiplier

$$\sqrt{\frac{r(\mathbf{a})}{r(\mathbf{b})}} \prod_{g\in\mathcal{B}} \prod_{i\in U_{\nu,g}} r_g(b_{g,i}) = \sqrt{\prod_{g\in\mathcal{B}} \prod_{i\in U_{\nu,g}} r_g(a_{g,i}) r_g(b_{g,i})}.$$

The shared-coordinate factors cancel exactly. The multiplier lies in $[1, \Lambda^t]$, with $t = (m+1)s_B$, so its difference from one has weighted row and reversed-column bounds at most $K(\Lambda^t - 1)$.

For the exclusion error, fix a source array. A free draw in group g can coincide with one of its M source coordinates or one of the $t_{\nu,g}$ newly integrated target coordinates. The remaining target coordinates are copied source coordinates. Under the harmonic laws, each of these events has probability at most $\kappa_g = \max_{p\in\mathcal{P}_g} \mu_g(p)$. The union bound gives

$$b_{\mathrm{col}} \geq \max_\nu \sum_{g\in\mathcal{B}} |C_{\nu,g}|(M + t_{\nu,g})\kappa_g.$$

The same calculation with source and target interchanged gives the reversed-column estimate. The tilt multiplier adds a factor at most Λ^t. Weighted Schur estimates therefore show, uniformly in Θ,

$$\left\| \mathcal{V}\widetilde{\mathcal{T}}_\zeta^{\mathrm{allow}}(\Theta)\mathcal{V}^{-1} - \mathcal{T}_\zeta(\Theta) \right\| \leq K(\Lambda^t - 1) + K\Lambda^t b_{\mathrm{col}}. \tag{5.52}$$

Consequently, the ideal estimate and (5.50)–(5.51) give

$$\|C\| \leq C_{\mathrm{cond}}\{\epsilon + K(\Lambda^t - 1) + K\Lambda^t b_{\mathrm{col}}\}.$$

The opposite orientation is the adjoint for the original mass. For a row kernel on $\mathcal{X}$, that adjoint has coefficient

$$K^*(x, y) = \frac{m_X(y)}{m_X(x)}\overline{K(y, x)}.$$

On the shared-coordinate support one has the exact active-mass identities

$$w_B(\mathbf{a})\nu_{U_\nu}(\mathbf{b}) = w_B(\mathbf{b})\nu_{U_\nu}(\mathbf{a}), \qquad D_{\nu,B}(\mathbf{a}, \mathbf{r}) = D_{\nu,B}(\mathbf{b}, \mathbf{r}).$$

The first follows by separating the shared and unshared factors in w_B. The second follows because every source label occurring in $D_{\nu,B}$ is shared. The free-exclusion indicator is symmetric in the two endpoint arrays. Slot permutations preserve w_B, and $\pi(T_h\mathfrak{m}) = \pi(\mathfrak{m})$. These equalities verify the weighted-adjoint formula term by term, including the original free law. Swapping endpoints conjugates the coefficient and phase, interchanges the two independent endpoint averages, and changes T_h to T_{-h}. Thus the reverse Fourier block is explicitly

$$\mathcal{B}^-_{\mathbf{d},\chi} = d^{-i\zeta} \sum_D e(-\theta kdD)\overline{\chi(kdD)}B_D^* = (\mathcal{B}^+_{\mathbf{d},\chi})^*.$$

It has the same norm, and the same completion and Parseval estimates therefore give the identical clean bound in both orientations.

For P, freeze the pattern, the free labels, the big part of the source permutation, and the entire target permutation. The remaining small source permutations choose one omitted slot in each small group, uniformly among $(J+1)^{s_S}$ choices. Let U be the product of the

omitted primes, and let S_S be the product of all small source labels. A pending particle (p, z), eligible at big dilation D_B, can reach residue zero only if

$$zU + (\pm k)D_B S_S = 0 \pmod{p}.$$

The second term is nonzero. Every omission product is below p, since

$$U \le \exp(s_S L^{36/100}) < \exp(L^{39/100}) \le p.$$

Unique factorization and the distinct small labels make these omission products distinct. Each eligible pending particle therefore permits at most one omission choice. A multiset of size at most B_0 permits at most B_0 choices. Newly stored particles have nonzero translated residue and supply no additional promotion choices.

The positive full edge and its zero-retrieval version each have the row and reversed column bounds just computed. Their difference costs a factor 2. Averaging the source-supported kernels gives a row bound multiplied by $B_0/(J+1)^{s_S}$, while preserving the column bound. Schur's test gives

$$\|P\| \le 2KC_{\mathrm{row}}C_{\mathrm{mem}}\sqrt{B_0/(J+1)^{s_S}}.$$

The mass-weighted adjoint swaps store and retrieval factors and reverses the edge orientation. Applying the same argument to it gives the identical bound for S_1. Adding the three bounds proves (5.35). □

For fixed m, C_1, (5.14) gives $\Lambda^t - 1 = L^{O(1)}e^{-L^{39/100}}$ and the same type of bound for b_{col}. Also

$$C_{\mathrm{cond}}, C_{\mathrm{row}} \le L^{1/5}, \qquad C_{\mathrm{mem}} \le L^{C_m}, \qquad C_m = \frac{7}{100}(m+1)\log\bigl(68(m+17)\bigr).$$

The last bound follows by substituting $\vartheta^{m+1} \ge 1/2$ and $\Lambda \le 2$ in the Schur constant. If $B_0 = \lceil L^2 \rceil$, the final term in (5.35) is at most

$$4\sqrt{2}\, L^{C_1+1/5+C_m+1-(9/400)\log(J+1)}.$$

Consequently, for every prescribed $G > 0$, a sufficiently large fixed J, chosen after m, C_1, G, gives

$$\|E\| \le L^{-G} \quad \text{if} \quad \sup_{\Theta} \|\mathcal{T}_\zeta(\Theta)\| \le L^{-(G+3)}. \tag{5.53}$$

For example it suffices to require

$$\frac{9}{400}\log(J+1) \ge G + \max(C_1, 0) + \frac{1}{5} + C_m + 4.$$

The constants governing the powers of L have already been fixed.

5.4. **The signed expansion and its boundary mass.** Use $B_0 = \lceil L^2 \rceil$, and let

$$\mathcal{I} = (\mathcal{B} \times \{1, \dots, M\}) \sqcup (\{0, \dots, n-1\} \times \mathcal{B} \times \{1, \dots, M\}) \sqcup (\{0, \dots, n\} \times \{0, \dots, B_0 - 1\}).$$

The three parts label initial births, fresh edge births in original slots, and ordered fresh ghost births. A history follows the order

$$G_0, E_0, G_1, E_1, \dots, E_{n-1}, G_n.$$

It records each full edge choice and its store mask, both endpoint permutations, the original free tuple, the complete target labels and every fresh or retrieval choice. Each ghost records all retention counts and its ordered fresh batch. Stored, retained and retrieved occurrences preserve their earlier birth. The corresponding addresses record precisely the fresh occurrences.

Assign a distinct symbol $\perp_a$ to every unused address a. For a history h, let $\widehat{p}_h(a)$ be its performed tagged prime, or $\perp_a$ when unused. Define

$$\Delta_\sigma(h) = \prod_{a\in\mathcal{I}} \mathbf{1}_{\widehat{p}_h(\sigma a)=\widehat{p}_h(a)}, \qquad r(\sigma) = |\mathcal{I}| - \#\{\text{cycles of } \sigma\}.$$

The cycle count includes the fixed points. The determinant identity gives

$$\sum_{\sigma\in\mathrm{Sym}(\mathcal{I})} \mathrm{sgn}(\sigma)\Delta_\sigma(h) = \mathbf{1}_{\text{all performed birth primes are distinct}}. \tag{5.54}$$

Indeed, the matrix with entries $\mathbf{1}_{\widehat{p}_h(a)=\widehat{p}_h(b)}$ is the identity when all values are distinct, and has two equal rows otherwise. The unused symbols contribute identity rows.

Let

$$Q_S = \prod_{g\in\mathcal{S}} \prod_{p\in\mathcal{P}_g} p, \qquad \tau_S = \prod_{g\in\mathcal{S}} V_g^{-M}.$$

For each $z \in [0, Q_S) \cap \mathbb{Z}$, let $W_\sigma(z)$ be the sum of the original history weights multiplied by Δ_σ, with initial coefficient

$$L_z(x) = \mathbf{1}_{\mathrm{position}(x)=z}\mathbf{1}_{\mathrm{pending}(x)=\varnothing} \prod_{p \text{ initially active}} \frac{m_p}{V_{g(p)}}$$

and terminal coefficient

$$R_z(x) = \mathbf{1}_{\mathrm{position}(x)=z}\mathbf{1}_{\mathrm{pending}(x)=\varnothing}.$$

The terminal active labels and the terminal set of previously used primes are summed freely. Each $W_\sigma(z)$ retains the operational zero guards and the factor Δ_σ. Global distinctness of performed births is supplied by (5.54) after the signed sum. Write

$$w_\sigma = \beta\frac{\tau_S}{Q_S} \sum_{0\le z<Q_S} W_\sigma(z), \qquad \mathfrak{p}_n = \sum_\sigma \mathrm{sgn}(\sigma) w_\sigma. \tag{5.55}$$

Equations (5.15) and (5.16), followed by (5.54), identify $\mathfrak{p}_n$ with the capped memory expansion of the original independent-residue moment.

Here is the boundary normalization in (5.55). For this calculation suppose $[0, Q_S)\cap\mathbb{Z} \subseteq \mathcal{W}$, as the window below will satisfy. Let C_z count the small physical label arrays at z, and put

$$S_0 = \sum_{\substack{\mathbf{p}_B \\ \text{distinct within each group}}} \prod_{g\in\mathcal{B}} \prod_{i=1}^{M} \nu_g(p_{g,i}).$$

For the row-kernel convention used here, the initial Hilbert vector is $\overline{L_z(x)/m_X(x)}$, where m_X is the state mass. At empty pending memory this is $\tau_S^{-1}R_z(x)$. The terminal Hilbert vector is R_z. The two boundary vectors therefore have squared norms $\tau_S^{-1}S_0C_z$ and $\tau_S S_0 C_z$, respectively. Their product is S_0C_z. Averaging the small divisibility conditions over their complete period gives

$$\frac{\tau_S}{Q_S} \sum_{0\le z<Q_S} C_z = \prod_{g\in\mathcal{S}} \sum_{\substack{p_1,\dots,p_M\in\mathcal{P}_g \\ \text{distinct}}} \prod_{i=1}^{M} \mu_g(p_i) \le 1. \tag{5.56}$$

Thus the normalized product of boundary norms is at most $\beta S_0 \le S_0$. Moreover, (5.14) gives

$$S_0 \le \left(1 + \frac{2(n+1)}{P_*}\right)^{Ms_B} \le 2 \tag{5.57}$$

for each fixed J and sufficiently large L.

Lemma 5.4 (Rank bounds)**.** *Fix m, C_1, and choose a constant $C_a \geq 0$, independent of J, large enough to dominate the absolute edge row bounds and the phase-modified edge bounds. Put $C_h = C_a + 8$. For each fixed J, sufficiently large L gives*

$$|w_\sigma| \leq 2L^{C_h(n+1)} \left(\frac{64}{P_*}\right)^{r(\sigma)}, \tag{5.58}$$

$$|w_\sigma| \leq L\, L^{(-G+8)n+8+2r(\sigma)(G+C_a)} \quad \textit{under } (5.53). \tag{5.59}$$

All thresholds are uniform in the finite coefficient family subject to its stated bound, in the phases, and in the permitted translations.

Proof. Put $T_0 = m + 1$, $v = 16$, and $\varrho = 1/2$. We use the nonnegative weight

$$\mathfrak{w}(x) = v^{|S(x)|},$$

where $S(x)$ is the pending multiset of the memory state x. The active labels and the small physical vertex remain part of x. All sums below retain the original finite choice spaces. In particular, a complete edge choice specifies both endpoint permutations, its pattern, its ordered free array, its target, its source-store mask, and each target-slot fresh or retrieval decision.

Chronological pivots and performed addresses. Fix an order on the big groups. Order the potential birth addresses by placing all initial addresses first, followed by

$$G_0, E_0, G_1, E_1, \ldots, E_{n-1}, G_n.$$

Within an edge, use the fixed group order and decreasing original unshared-slot order. Within a ghost, use increasing position in its ordered fresh batch. This is the order in which the corresponding fresh choices are made.

In every nontrivial cycle C of σ, choose its least address a_C. Define

$$\mathcal{D}_\sigma = \bigcup_{|C|>1} (C \setminus \{a_C\}), \qquad \mathrm{piv}(a) = a_C \quad (a \in C \setminus \{a_C\}).$$

Then

$$|\mathcal{D}_\sigma| = \sum_C (|C| - 1) = r(\sigma), \qquad \mathrm{piv}(a) < a \quad (a \in \mathcal{D}_\sigma). \tag{5.60}$$

On the event $\Delta_\sigma = 1$, every moved address is performed. Indeed, the value of an unused address is its own symbol $\perp_a$, which occurs at that address alone. The performed tagged prime is constant on each nontrivial cycle.

The initial labels are distinct within every group, and the prime groups are disjoint. Their tagged values are consequently all distinct. If a cycle contains two initial addresses, its contribution to $W_\sigma(z)$ is zero. For every remaining σ, a cycle containing an initial address has that address as its pivot. Thus every member of $\mathcal{D}_\sigma$ belongs to a fresh edge or ghost birth. Write

$$r_i^E = |\mathcal{D}_\sigma \cap \mathcal{I}(E_i)|, \qquad r_j^G = |\mathcal{D}_\sigma \cap \mathcal{I}(G_j)|.$$

These integers depend on σ and the potential address sets, and satisfy

$$\sum_{i=0}^{n-1} r_i^E + \sum_{j=0}^{n} r_j^G = r(\sigma). \tag{5.61}$$

A record $\mathcal{R}$ stores the tagged prime at each earlier performed address. Initialize it with all the initial labels. After a fresh draw at address a, first test the required equality and then write its value at a. All earlier entries remain available through stores, translations, retentions, retrievals and deaths.

For an operation o, its conditional guard $\mathcal{G}_o$ requires every address in $\mathcal{D}_\sigma \cap \mathcal{I}(o)$ to be performed. At each such fresh address a, it also requires

$$p(a) = \mathcal{R}(\mathrm{piv}(a)), \tag{5.62}$$

using the record immediately before that draw. The test fails if its pivot entry is unavailable or belongs to a different group. On a history with nonzero original coefficient, induction through the chronological order shows that $\Delta_\sigma = 1$ implies all these conditional guards. Histories with zero coefficient contribute zero to the estimate.

The performed requirement fixes the number of prescribed draws in every contributing local sum. For an edge, a dependent original slot must lie in that pattern's unshared set and must use the fresh alternative. For a ghost of length b, every dependent ghost address must have index less than b. In either case an omitted dependent address makes the conditional contribution zero.

One target-slot sum. For a big group g, let

$$\mathcal{Z}_g = \{(p, r) : p \in \mathcal{P}_g,\ r \in \mathbb{Z}/p\mathbb{Z}\}, \qquad z_g(S) = \sum_{p \in \mathcal{P}_g} S(p, 0).$$

In this one-group calculation, S denotes the component of the pending multiset in $\mathcal{Z}_g$. The exact fresh weight is $\nu_g(p) = m_p/V_g$. By (5.14),

$$\sum_{p \in \mathcal{P}_g} \nu_g(p) \le 2, \qquad \nu_g(p) \le \frac{4}{P_*}. \tag{5.63}$$

For a fixed displacement, the store and retrieval factors satisfy

$$\beta_s(p) = \mathbf{1}_{p \nmid D_B}, \qquad \beta_r(p) = V_g^{-1}\mathbf{1}_{p \nmid D_B}, \qquad 0 \le \beta_s(p), \beta_r(p) \le 2.$$

Set

$$\vartheta = 1 - \frac{1}{2(T_0 + 1)}, \qquad \Phi_k(S) = \vartheta^{k z_g(S)} v^{|S|}, \qquad 0 \le k \le T_0.$$

Bernoulli's inequality gives $\vartheta^{T_0} \ge 1 - T_0/(2(T_0 + 1)) > 1/2 = \varrho$.

At an unprescribed target slot, the exact two alternatives, applied to a nonnegative future function F, are

$$\sum_{p \in \mathcal{P}_g} \left\{ \nu_g(p) F(S, \mathcal{R}[a \mapsto p]) + \beta_r(p) S(p, 0) F(S - [(p, 0)], \mathcal{R}) \right\}. \tag{5.64}$$

Terms with $S(p, 0) = 0$ in the second summand are zero. The multiplicity $S(p, 0)$ is the original retrieval multiplicity. At a prescribed slot, the sum consists of the fresh term with the value in (5.62).

The fresh part of (5.64), with future weight Φ_{k+1}, is at most $2\Phi_k(S)$. Each one-slot application has $0 \le k < T_0$. For the retrieval part, write $h = z_g(S)$. If $h \ge 1$, then

$$\begin{aligned} &\sum_p \beta_r(p) S(p, 0) \Phi_{k+1}(S - [(p, 0)]) \\ &\qquad \le 2h\vartheta^{(k+1)(h-1)} v^{|S|-1} \le \frac{2\vartheta^{-k}}{v(1 - \vartheta)} \Phi_k(S). \end{aligned} \tag{5.65}$$

For the last inequality, use

$$h\vartheta^{h-1} \le \sum_{a=0}^{h-1} \vartheta^a \le \frac{1}{1 - \vartheta}.$$

When $h = 0$, the retrieval sum is zero. Hence every unprescribed slot costs at most

$$H_{T_0} = 2 + \frac{2\vartheta^{-T_0}}{v(1 - \vartheta)}. \tag{5.66}$$

A prescribed slot costs at most

$$\nu_g(\mathcal{R}(\mathrm{piv}(a)))\,\Phi_{k+1}(S) \le \delta\Phi_k(S), \qquad \delta = \frac{64}{P_*},$$

by (5.63). These inequalities also apply when the future function depends on all the active labels and on the entire updated record. Namely, first bound that future function uniformly by a constant times Φ_{k+1}, and then apply the displayed sums. The record in a prescription is fixed before the corresponding fresh sum is taken.

The complete conditional edge row. Fix the chronological source, the two endpoint permutations, the pattern, and the full free tuple. The displacement uses only the shared labels and that free tuple. Retain the shared-label and position guards at this stage. On a nonzero original scalar they identify the shared labels at the two endpoints. In particular, a reversed edge has the same source-determined displacement after interchanging its endpoint permutations and replacing k by $-k$. Thus both the translation by the full displacement and the eligibility product D_B are fixed before any target fresh draw.

For the chosen pattern, list its unshared original slots increasingly and write $t_g \le T_0$ for their number. A source-store mask $\alpha = (\alpha_{g,a})$ inserts the indicated old source labels at residue zero. If S^α is the resulting multiset, then

$$|S^\alpha| = |S| + \sum_{g,a} \alpha_{g,a}.$$

Translate this multiset by the full displacement and process the target slots in decreasing original order, group by group. Translation preserves its cardinality. The terminal damping of group g satisfies

$$\varrho^{z_g(S')} v^{|S'|} \le \Phi_{t_g}(S').$$

Pulling this inequality backwards through the t_g exact slot sums, using (5.65), gives

$$v^{|S^\alpha|} \prod_g \prod_{a=1}^{t_g} \begin{cases} \delta, & \text{the original address is dependent,} \\ H_{T_0}, & \text{otherwise.} \end{cases} \tag{5.67}$$

The source damping is at most one. Each intermediate cutoff and each original zero coefficient is still included in the local choice sum. Enlarging the nonnegative sum by removing a cutoff preserves this upper bound.

Summing the original source-store masks gives the exact identity

$$\begin{aligned} \sum_\alpha \left(\prod_{g,a} \beta_s(p^{\mathrm{src}}_{g,a})^{\alpha_{g,a}} \right) v^{|S| + \sum_{g,a} \alpha_{g,a}} \\ = v^{|S|} \prod_{g,a} (1 + v\beta_s(p^{\mathrm{src}}_{g,a})) \le v^{|S|} \prod_g (1+2v)^{t_g}. \end{aligned} \tag{5.68}$$

Here the product convention gives factor one when a slot is not stored. Repeated particles in the pending multiset keep their full multiplicities throughout (5.65) and (5.68).

If the conditional guard is possible for this pattern, the number of prescribed positions in (5.67) equals the fixed owner count r_i^E. Consequently the memory part of the edge row is at most

$$\delta^{r_i^E} v^{|S|} K_{T_0}^{s_B}, \qquad K_{T_0} = \left\lceil 33 \max\left\{1, 2 + \frac{2\vartheta^{-T_0}}{16(1-\vartheta)}\right\} \right\rceil^{T_0}. \tag{5.69}$$

If a dependent original address is absent from the pattern, or its fresh alternative is absent, the guarded contribution is zero.

It remains to sum the physical data fixed at the start of this calculation. A complete memory trace fixes the selected target array. For fixed endpoint permutations, the map from

a compatible target to its small physical vertex and selected target array is injective. Thus the remaining target sum is bounded by the small-label completion sum at the source-determined target position.

In a small group there is one unshared slot. If h eligible primes divide the target position outside the fixed shared labels, the sum of its target damping and normalization is at most

$$\frac{h(\sqrt{q})^{h-1}}{V_g} \le \frac{1}{(1-\sqrt{q})V_g} \le \frac{32}{15}.$$

For $h=0$ the sum is zero. We may therefore use the uniform bound $C_S=(47/15)^{s_S}$, in agreement with the sum of the zeroth and first target-completion terms.

Bound the original coefficient by $K_i(\nu)$. The phase has absolute value one. After the source displacement has been fixed as above, removing the free-label exclusions enlarges a nonnegative sum. The original free law and both permutation averages satisfy

$$\sum_{\pi_s,\pi_t} \frac{1}{(M!)^{2s}} \sum_{\nu} K_i(\nu) \sum_{\mathbf{r}} \prod_{g} \prod_{a\in C_{\nu,g}} \mu_g(r_{g,a}) = \sum_{\nu} K_i(\nu). \tag{5.70}$$

The free sum in this formula uses its original ordered space for each pattern. Combining (5.69)– (5.70) proves, for every source x and every incoming record $\mathcal{R}$,

$$\sum_{c\in\mathcal{C}_i(x)} \mathcal{G}_{E_i}(x,\mathcal{R},c)\,|a_i(x,c)|\,\mathfrak{w}(x_i') \le E_i \delta^{r_i^E}\,\mathfrak{w}(x), \qquad E_i=\left(\sum_{\nu} K_i(\nu)\right) C_S K_{T_0}^{s_B}. \tag{5.71}$$

Here a_i is the full original edge coefficient, including all its operational zero guards, and x_i' is its original next state. The estimate is uniform in the target-dependent coefficient and in the incoming active labels and record.

The complete conditional ghost row. Put $\mathcal{Z}=\bigsqcup_{g\in\mathcal{B}}\mathcal{Z}_g$. Write $u=(p,r)$, and let s_u be the incoming multiplicity of u. For retained counts $0\le k_u\le s_u$, put

$$\mathcal{B}_j(s,k)=\prod_u \binom{s_u}{k_u} A_j(u)^{k_u} D_j(u)^{s_u-k_u}.$$

For a nonnegative future function F, the original ordered ghost action is

$$\sum_{0\le k\le s} \mathcal{B}_j(s,k) \sum_{b=0}^{B_0-|k|} \frac{1}{b!} \sum_{(u_0,\dots,u_{b-1})\in\mathcal{Z}^b} \left(\prod_{a=0}^{b-1} C_j(u_a)\nu_{g(u_a)}(p(u_a))\right) F\left(k+\sum_{a<b}[u_a]\right). \tag{5.72}$$

The other physical coordinates are preserved. The sum includes repetitions. A fresh histogram (h_u), of total size b, has $b!/\prod_u h_u!$ ordered representatives. Thus the factor $1/b!$ in (5.72) gives precisely its original multiset coefficient $\prod_u 1/h_u!$.

Let $D_j^{\mathrm{addr}}\subseteq\{0,\dots,B_0-1\}$ be the dependent ghost positions. The performed guard requires $D_j^{\mathrm{addr}}\subseteq\{0,\dots,b-1\}$, so every admissible batch has exactly r_j^G prescribed coordinates. After including the output factor v for each fresh particle, its positive fresh weight is

$$\omega_j(p,r)=v|C_j(p,r)|\nu_{g(p)}(p)=\mathbf{1}_{r=0}\frac{v}{\varrho}\eta_j m_p. \tag{5.73}$$

In particular,

$$\sum_{r \bmod p} \omega_j(p,r)=\frac{v}{\varrho}\eta_j m_p\le\delta, \qquad \sum_{p,r}\omega_j(p,r)\le\Lambda_G, \qquad \Lambda_G=32\sum_p m_p.$$

The factor V_g cancels in (5.73), and there is one nonzero residue contribution for each tagged prime.

For fixed retained counts and batch length, integrate the ordered fresh coordinates backwards, conditioning each time on the entire earlier prefix and its record. A free coordinate costs at most Λ_G. A prescribed coordinate costs at most δ. The resulting bound is

$$\sum_{u_0,\dots,u_{b-1}} \mathbf{1}_{\text{record tests}} \prod_{a<b} \omega_j(u_a) \le \delta^{r_j^G} \Lambda_G^{b-r_j^G}.$$

It holds uniformly in the initial record, and therefore also when a pivot belongs to an earlier position of the same batch. Keep the original factorial and then enlarge the permitted lengths to obtain

$$\sum_{b\ge r_j^G} \frac{\delta^{r_j^G} \Lambda_G^{b-r_j^G}}{b!} = \delta^{r_j^G} \sum_{h\ge 0} \frac{\Lambda_G^h}{(h+r_j^G)!} \le \delta^{r_j^G} e_G^{\Lambda}. \tag{5.74}$$

All prescribed positions are fixed potential addresses in this sum.

Finally, the retained-particle sum is the binomial identity

$$\sum_{0\le k\le s} |\mathcal{B}_j(s,k)| v^{|k|} = \prod_u \left(A_j(u)v + |D_j(u)|\right)^{s_u} \le v^{|s|}. \tag{5.75}$$

For nonzero residues the factor in parentheses is v. At zero, $1-\eta_j \le 1/16$, and hence

$$A_j(p,0)v + |D_j(p,0)| = \frac{1-\eta_j}{\varrho^2} v + \frac{\eta_j}{\varrho} \le \frac{1/16}{(1/2)^2}\, 16 + 2 = 6 \le 16.$$

Equations (5.74) and (5.75), applied first with the common cutoff in (5.72), give

$$\sum_{c\in\mathcal{C}_{G_j}(x)} \mathcal{G}_{G_j}(x,\mathcal{R},c)\, |a_{G_j}(x,c)|\, \mathfrak{w}(x_j') \le G_+ \delta^{r_j^G} \mathfrak{w}(x), \qquad G_+ = \exp\left(32 \sum_p m_p\right). \tag{5.76}$$

Conditional multiplication along the original histories. List the $2n+1$ operations in their chronological order. For operation o, denote its full original coefficient by $a_o(x,c)$, its next state by $x_o'(x,c)$, and its deterministic record update by $U_o(\mathcal{R},c)$. The coefficient includes its original cutoff, compatibility and eligibility guards. For a fixed root z, define recursively

$$F_{2n+1}(x,\mathcal{R}) = |R_z(x)|,$$

$$F_o(x,\mathcal{R}) = \sum_{c\in\mathcal{C}_o(x)} |a_o(x,c)| \mathcal{G}_o(x,\mathcal{R},c) F_{o+1}(x_o'(x,c), U_o(\mathcal{R},c)).$$

These are finite sums over the original source-dependent choices. The terminal bound is $F_{2n+1} \le \mathfrak{w}$. If the future function satisfies $F_{o+1}(x,\mathcal{R}) \le C\mathfrak{w}(x)$ for every state and every record, then (5.71) or (5.76) gives the same bound at operation o, multiplied by its local constant and prescribed-draw factor. This induction applies with all correlations in the target choices and in the record present.

Along a history with nonzero coefficient and $\Delta_\sigma = 1$, all conditional guards equal one. Taking absolute values in the original finite history sum and using (5.61) therefore gives

$$|W_\sigma(z)| \le \left(\prod_{i=0}^{n-1} E_i\right) G_+^{n+1} \delta^{r(\sigma)} \sum_x |L_z(x)|. \tag{5.77}$$

The record is a deterministic function of the exposed prefix, so the induction introduces no sum over possible records. The case of a dependent initial address already has value zero. The initial pending multiset is empty, so $\mathfrak{w}(x) = 1$ on the support of L_z. This gives the initial sum in (5.77).

By the exact boundary normalization in (5.55) and (5.56),

$$\beta\frac{\tau_S}{Q_S}\sum_{0\le z<Q_S}\sum_{x}|L_z(x)|\le S_0.$$

In particular, the original lifespan baseline β, the small physical mass τ_S, and the inverse period Q_S^{-1} have all been included. Summing (5.77) gives

$$|w_\sigma|\le S_0\left(\prod_{i=0}^{n-1}E_i\right)G_+^{n+1}\left(\frac{64}{P_*}\right)^{r(\sigma)}. \tag{5.78}$$

The uniform polynomial factor. The complete coefficient budget is $\sum_\nu K_i(\nu)\le 1+L^{C_1}$. Define

$$C_e=\max(0,C_1)+1+\frac{9}{100}\log(47/15)+\frac{7}{100}\log K_{T_0}.$$

The group counts in (5.1) and (5.71) imply $E_i\le L^{C_e}$ for $L\ge 2$. Also

$$\sum_p m_p\le 2\sum_{g\in\mathcal{B}}V_g\le 3s_B,\qquad G_+\le\exp(96s_B)\le L^7.$$

For fixed J, the full initial mass satisfies

$$S_0\le\left(1+\frac{2(n+1)}{P_*}\right)^{Ms_B}\le\exp\left(\frac{2Ms_B(n+1)}{P_*}\right)\le 2$$

once L is sufficiently large. The last assertion uses $n=2\lceil L^{52/100}\rceil$, $s_B=O(\log L)$, and $P_*=\exp(L^{39/100})$. Choose $C_a\ge C_e$ and put $C_h=C_a+8$. Equation (5.78) now yields

$$|w_\sigma|\le 2L^{C_en+7(n+1)}\left(\frac{64}{P_*}\right)^{r(\sigma)}\le 2L^{C_h(n+1)}\left(\frac{64}{P_*}\right)^{r(\sigma)}.$$

The exponents C_e,C_a,C_h depend on the fixed replacement bound and coefficient budget. They are independent of J. The final threshold is allowed to depend on the fixed J. Every displayed local estimate is uniform in the original phases, translations, target choices and incoming record. This proves (5.58).

For the second estimate, choose a sign function ϵ_a on the finite alphabet of tagged primes and unused symbols for every potential address a. All signs $\epsilon_a(u)$, indexed by the pairs (a,u), are independent and uniform in $\{-1,1\}$. For two values u,v, $\mathbb{E}\epsilon_a(u)\epsilon_a(v)=\mathbf{1}_{u=v}$. Consequently Δ_σ is the average of a product of local phases of absolute value one. At address a, the local factor is $\epsilon_a(\widehat{p}_h(a))\epsilon_{\sigma^{-1}a}(\widehat{p}_h(a))$. It is 1 at every fixed address. The support of σ has size at most $2r(\sigma)$, so at most $2r(\sigma)$ edge operations acquire a phase. The original signed sum over edge choices is retained at every unaffected operation.

An unchanged edge has norm at most L^{-G}. For a changed edge, the weighted row and reversed column Schur estimates give norm at most L^{C_a}, even for these choice-dependent phases. The phases on initial addresses multiply the initial boundary vector by a function of modulus one.

We also give the two-sided ghost estimate. Fix the active arrays and the small physical vertex. For $u=(p,r)$, write $\lambda_u=\nu_{g(p)}(p)$. The absolute ghost kernel is bounded by the product of the one-coordinate kernels

$$H_{j,u}(a,b)=\sum_{h=0}^{\min(a,b)}\binom{a}{h}A_j(u)^h|D_j(u)|^{a-h}\frac{(|C_j(u)|\lambda_u)^{b-h}}{(b-h)!}.$$

Expanding the ordered births and collecting equal entries gives exactly this expression. Arbitrary local phases have modulus one and obey the same absolute bound. For the weight

$v^{|s|}$, the unrestricted row sum is

$$\exp\left(v\sum_u |C_j(u)|\lambda_u\right)\prod_u \left(A_j(u)v+|D_j(u)|\right)^{s(u)}.$$

The exact factorial balance is

$$\pi(s)H_{A,|D|,|C|}(s,t)=\pi(t)H_{A,|C|,|D|}(t,s).$$

It follows by writing the term with retained multiplicity h as

$$\frac{\lambda_u^{a+b-h}A_j(u)^h|D_j(u)|^{a-h}|C_j(u)|^{b-h}}{h!(a-h)!(b-h)!}$$

in each coordinate. The reversed-column sum therefore has the same row formula with C_j and D_j interchanged. At zero residue,

$$A_j=4(1-\eta_j)\le\frac{1}{4},\qquad |D_j|=2\eta_j\le 2,\qquad |C_j|=2\eta_j V_g\le 4.$$

At a nonzero residue $A_j=1$ and $C_j=D_j=0$. Both $A_j+|D_j|/16\le 1$ and $A_j+|C_j|/16\le 1$ hold. The weighted row and column constants can thus be taken to be

$$R_j=\exp\left(16\sum_u|C_j(u)|\lambda_u\right),\qquad C_j^{\mathrm{Schur}}=\exp\left(16\sum_u|D_j(u)|\lambda_u\right).$$

Deleting terms to impose the common total cutoff preserves these positive bounds. The fixed active-array and small-vertex factor in the state mass cancels in the reversed-column calculation. Schur's inequality gives, for each ghost with any of the phases,

$$\|G_j\|\le\sqrt{R_jC_j^{\mathrm{Schur}}}\le\exp\left(48\sum_{g\in\mathcal{B}}\sum_{p\in\mathcal{P}_g}\nu_g(p)\right)\le e^{96s_B}\le L^8.$$

Only the zero residue contributes to the sums in the exponent.

Apply Cauchy–Schwarz to the boundary vectors and multiply these operator norms in chronological order. Equations (5.56) and (5.57) give

$$|w_\sigma|\le 2L^{8(n+1)}L^{-Gn}\left(L^{G+C_a}\right)^{2r(\sigma)}.$$

Since $2\le L$, this is (5.59).

The phase-modified Schur exponent is bounded by

$$C_\ell=\max(0,C_1)+1+\frac{1}{5}+\frac{7}{100}T_0\log\bigl(68(T_0+16)\bigr).$$

Taking $C_a=\max(0,C_e,C_\ell)$ establishes the asserted independence from J. □

5.5. **Summation of ranks and transfer to the physical graph.** Continue to use these values of R and n, and set

$$B_0=\lceil L^2\rceil,\qquad U=\lfloor e^L\rfloor,$$

$$H=nU^{\,2Ms+1},\qquad \mathcal{W}=[-H,-H+2Q_S)\cap\mathbb{Z}. \tag{5.79}$$

For fixed J, (5.1) implies

$$|\mathcal{I}|=Ms_B+nMs_B+(n+1)B_0\le L^4$$

eventually. The first big group has at least $P_*/2$ primes, so the capacity B_0 is available. The first small group likewise gives $\log Q_S\ge\frac{1}{2}e^{L^{27/100}}L^{27/100}$. It follows that

$$2H\le Q_S,\qquad 2X\le Q_S\quad\text{whenever}\quad e^{3L/4}\le X\le e^{5L/4}. \tag{5.80}$$

If $|k| \le U$, every path of n edges starting in $[X, 2X)$ remains in $\mathcal{W}$, because each displacement has absolute value at most U^{2Ms+1}.

The number of permutations of $\mathcal{I}$ having rank r is at most $|\mathcal{I}|^{2r}$. One proof chooses a transposition decomposition with r factors by opening the cycles. There are at most $|\mathcal{I}|^2$ choices for each transposition. Fix $E > 0$, take $G = E + 9$, and put

$$r_0 = \left\lceil \frac{2(C_h + E + 5)(n+1)\log L}{L^{39/100}} \right\rceil. \tag{5.81}$$

For large L,

$$|\mathcal{I}|^2 \frac{64}{P_*} \le 256 L^8 e^{-L^{39/100}} \le e^{-L^{39/100}/2} \le \frac{1}{2}.$$

The geometric sum and (5.58) therefore give

$$\sum_{r(\sigma) \ge r_0} |w_\sigma| \le 4 L^{C_h(n+1)} e^{-r_0 L^{39/100}/2} \le L^{-En}.$$

For $r < r_0$, (5.59) gives

$$\sum_{r(\sigma) < r_0} |w_\sigma| \le r_0 L^{(-G+8)n + 9 + \{8 + 2(G + C_a)\} r_0} \le L^{-En}.$$

The last inequality uses $r_0/n \longrightarrow 0$ and $G - 8 = E + 1$. Thus

$$|\mathfrak{p}_n| \le 2L^{-En}. \tag{5.82}$$

We state the two approximation bounds used at this final step. Let $\mathfrak{m}_n$ be the independent-residue average of the original alternating path expansion. Uniformly for $|k| \le U$, $p \nmid k$ for every group prime, and $e^{3L/4} \le X \le e^{5L/4}$, every fixed $E > 0$ gives

$$\left| \sum_{z \in [X, 2X) \cap \mathbb{Z}} \left\langle u_z, (\mathcal{A}_k^{\zeta,\theta} (\mathcal{A}_k^{\zeta,\theta})^*)^R u_z \right\rangle - X \mathfrak{m}_n \right| \le X L^{-En}, \tag{5.83}$$

$$|\mathfrak{m}_n - \mathfrak{p}_n| \le L^{-En}. \tag{5.84}$$

For (5.83), fix all path choices. The active congruences have modulus with logarithm $O_J(nTL^{46/100}) = o(L)$. Only compatible active congruences contribute. Factor out their deterministic damping factors. The remaining independent-residue mean is at least $q_{\mathrm{act}}^{-1} \exp(-CnT)$, where q_{act} is the active modulus. Each remaining prime has hit probability at most $(n+1)/p = o(1)$, and the sum of these probabilities is $O(nT)$. Expand the remaining damping products by two adjacent Bonferroni truncations at order $C_E nT$. Their gap is bounded by the corresponding exponential-series tail, since the total reciprocal-prime intensity is $O(nT)$. Increasing C_E gives relative error L^{-En}. Every retained term is a CRT condition with modulus and total number of terms $\exp(O_E(nTL^{46/100})) = \exp(o(L))$. Counting such a condition in an interval of length X has absolute error at most one. The resulting normalized error is $\exp(-3L/4 + o(L))$. The absolute row majorants sum the path choices at cost $L^{O(n)}$, which this error absorbs. Coprimality makes multiplication by k bijective in every prime residue space. The buffer (5.80) supplies the stated position window for all these paths.

For (5.84), apply the positive doubled row estimates with the weight $16^{|\mathrm{pending}|}$. The visit factors satisfy $1 - \eta_j \le 1/16$. A history exceeding B_0 particles has at least $B_0 + 1 - nMs$ excess births after its possible edge stores. Since every group prime is at most U, the doubled absolute edge row is bounded by

$$L^{C_1} (2048 U^4)^{Ms}.$$

The ordered ghost factorials bound the remaining birth sums by $\exp(192s)$ per visit. With the same boundary calculation (5.56) and $S_0 \le 2$, the total discarded normalized mass is therefore at most

$$2 \cdot 2^{-B_0-1+nMs} \exp\bigl(192(n+1)s\bigr) \left\{ L^{C_1} (2048U^4)^{Ms} \right\}^n .$$

In particular it is bounded by

$$2^{-B_0-1+nMs} \exp\bigl(O_{m,C_1,J}(nLT)\bigr).$$

Since $B_0 \asymp L^2$, $nLT = O(L^{38/25} \log L) = o(L^2)$, and $nMs = O_J(L^{13/25} \log L)$, this bound is at most L^{-En} for every fixed E.

Theorem 5.5 (Physical moment transference)**.** *Fix $A > 0$, $m \ge 1$ and $C_1 \ge 0$. There is $J_0 \ge m$ such that for every fixed $J \ge J_0$, and then for all sufficiently large L, the following holds. Let a finite pattern family satisfy*

$$C_{\nu,g} = \varnothing \ (g \in \mathcal{S}), \qquad |C_{\nu,g}| \le m \ (g \in \mathcal{B}), \qquad \sum_{\nu} \|c_\nu\|_\infty \le L^{C_1}.$$

If

$$\sup_{\Theta \in \mathbb{R}} \|\mathcal{T}_\zeta(\Theta)\| \le L^{-(A+13)},$$

then, with $\mathcal{W}$ as in (5.79), for every $\theta \in \mathbb{R}$, every integer k satisfying

$$|k| \le \lfloor e^L \rfloor, \qquad p \nmid k \quad (p \in \textstyle\bigcup_g \mathcal{P}_g),$$

and every natural X with $e^{3L/4} \le X \le e^{5L/4}$,

$$\sum_{z \in [X,2X) \cap \mathbb{Z}} \left\langle u_z, (\mathcal{A}_k^{\zeta,\theta} (\mathcal{A}_k^{\zeta,\theta})^*)^R u_z \right\rangle \le X \bigl(L^{-A}\bigr)^{2R}. \tag{5.85}$$

Proof. Take $E = A + 1$ and $G = E + 9 = A + 10$. Choose J_0 by (5.53), using ideal precision $G + 3 = A + 13$. Lemma 5.4 and (5.81) give (5.82). Combining (5.83), (5.84) and (5.82) bounds the left side of (5.85) by $4XL^{-(A+1)n}$. Since $n = 2R \ge 2$ and $L \ge 4$, this is at most XL^{-An}.

The alternating path expansion here closes the integer position. Its initial and terminal label arrays are summed independently. Expanding the row action of $(\mathcal{A}_k^{\zeta,\theta}(\mathcal{A}_k^{\zeta,\theta})^*)^R$ gives exactly these paths, with the initial vertex mass τ once. This proves that the quantity just bounded is the displayed physical root moment. □

Theorem 5.6 (The constructed physical moment)**.** *For every $A > 0$, there exist $m \ge 1$ and positive q_*, h_*, x_*, C_1, followed by an integer $J_0 \ge m$, with the following property. For every fixed $J \ge J_0$ and all sufficiently large L, there is a partition $\mathcal{B} = \mathcal{B}_1 \sqcup \mathcal{B}_2$ and the actual coefficient family (5.4) such that (5.85) holds simultaneously for*

$$|\zeta| \le L^2, \quad \theta \in \mathbb{R}, \quad |k| \le \lfloor e^L \rfloor, \quad p \nmid k \ (p \in \textstyle\bigcup_g \mathcal{P}_g), \quad e^{3L/4} \le X \le e^{5L/4}, \quad X \in \mathbb{N}.$$

The family is chosen before ζ, θ, k, X.

Proof. Apply Theorem 5.2 with $D = A + 13$. Its m, q_*, h_*, x_*, C_1 are fixed before J. Choose J_0 large enough for Theorem 5.5 and for $J_0 \ge m$. For each fixed $J \ge J_0$, choose the threshold in L after these fixed parameters. The partition and the coefficients then satisfy the complete ideal estimate for all the displayed frequencies. Theorem 5.5 gives the conclusion. □

For an integer $N > 1$, the substitution $L = \log N$ gives $\lfloor e^L \rfloor = N$. Hence $k = -N$ is included whenever the group primes are coprime to N. The residue changes and the position buffer used above apply to this value of k, with the same family and the same moment bound.

6. Reciprocal character energy

In this section the exponent of the center cutoff from Section 4 is denoted by d_0. Thus

$$L = \log N, \qquad Q = \exp(L^{d_0}), \qquad y = \exp(\varepsilon L), \qquad 0 < \varepsilon < \tfrac{1}{4}, \quad 0 < d_0 < \tfrac{1}{2}.$$

This notation permits d to retain its meaning as an endpoint of a prime-group range. Fix $0 < \kappa < d_0$ and a prime interval

$$\mathcal{P}_0 = (a, b] \cap \mathbb{P}, \qquad e^{L^\kappa} \le a < b \le Q, \qquad V = \sum_{p \in \mathcal{P}_0} \frac{1}{p} \ge \frac{1}{2}.$$

Let $r \ge 1$ satisfy $r \le L^D$, and suppose that N has no prime divisor at most $e^{\eta L}$, where D and $\eta > 0$ are fixed. All powers of L below have fixed exponents. Every auxiliary finite set and every coefficient array may depend on L.

Write $F(n)$ for the invariant center of Section 4. For sufficiently large N it satisfies

$$|F(n)| \le 1, \qquad F(vw) = F(v) \quad \text{if every prime divisor of } w \text{ is at most } Q. \tag{6.1}$$

The other marked groups are disjoint from $\mathcal{P}_0$ and contain only primes at most Q. Their coupled coefficient is denoted by $O(n)$. With damping $0 \le \vartheta \le 1/2$, at most T marks in each group, group masses at least $1/2$, at most $2 \log L$ groups, and label coefficients of absolute value at most L^E, the factorial marked-weight bound gives

$$|O(n)| \le L^{C_0}, \qquad C_0 = E + 2 \log(T! 4^T). \tag{6.2}$$

Indeed, the contribution of a group with $m \le T$ marks is at most

$$\frac{m!}{(1-\vartheta)^m V_g^m} \le T! 4^T.$$

Summing the coupled label coefficients once and multiplying these group bounds proves (6.2). Both C_0 and T are independent of the number of prime groups.

6.1. **Supported parts and the reduced modulus.** Define

$$u_r(n) = \prod_{p \mid r} p^{v_p(n)}.$$

Every positive integer has the unique factorization $n = uv$ in which $u = u_r(n)$ and $(v, r) = 1$. If $(n, N) = 1$ and $(N, r) = 1$, this becomes $(v, Nr) = 1$. Put

$$g = (u, r), \qquad q = N\frac{r}{g}, \qquad qg = Nr. \tag{6.3}$$

For u supported on the primes dividing r, the integer u/g is coprime to q. Its residue υ is therefore a unit modulo q. For every integer m and every h modulo qg,

$$e\left(\frac{umh}{qg}\right) = e\left(\frac{\upsilon mh}{q}\right). \tag{6.4}$$

The expression on the right depends only on h mod q, so a sum of its squared absolute value over h mod qg is exactly g times the corresponding sum over h mod q.

For completeness, the coprimality assertion follows prime by prime. A prime dividing both u/g and r/g would occur in u and r to exponents both exceeding their minimum, which is impossible. Also $(u, N) = 1$, since every prime divisor of u divides r and $(N, r) = 1$.

6.2. **The selected-prime arrays.** Let $\psi \in C_c^\infty((0,\infty))$ have support in $[1/2, 2]$ and $|\psi| \le 1$. We use the reciprocal physical weight

$$W_{\beta,\zeta}(n) = \frac{1}{n}\psi(n/Z)e(\beta n)n^{i\zeta} \qquad (n \ge 1). \tag{6.5}$$

It satisfies $|W_{\beta,\zeta}(n)| \le 2/Z$. Set $H_0 = \lfloor 8Z \rfloor$ and use the clipped dyadic intervals

$$a_j = \max(a, 2^j), \qquad b_j = \min(b, 2^{j+1}), \qquad \mathcal{P}_j = (a_j, b_j] \cap \mathbb{P} \quad (0 \le j < K),$$

where $b \le 2^K$. These intervals partition $\mathcal{P}_0$ with their open-left, closed-right endpoints retained, and $b_j \le 2a_j$. For a positive supported part u, put

$$\ell_j = \left\lfloor \frac{\lfloor H_0/u \rfloor}{\lfloor a_j \rfloor + 1} \right\rfloor, \qquad \mathcal{V}_j = \{1 \le v \le \ell_j : (v, Nr) = 1\}. \tag{6.6}$$

The same set $\mathcal{V}_j$ is used for every prime in $\mathcal{P}_j$. It contains every positive unit v for which $\psi(upv/Z) \ne 0$. In fact $upv \le 2Z$, while $p \le b_j \le 2a_j$ gives

$$(\lfloor a_j \rfloor + 1)uv \le 2puv \le 4Z \le H_0 \qquad (Z \ge 1).$$

Additional terms introduced by (6.6) have zero physical weight. Also

$$\ell_j \le \frac{8Z}{ua_j}. \tag{6.7}$$

We use the physical height condition

$$p^2 uNr \le H_0 \qquad (p \in \mathcal{P}_0). \tag{6.8}$$

For a nonempty $\mathcal{P}_j$, it implies $a_j uq \le 8Z$. Thus both q and ℓ_j are at most $8Z/(ua_j)$.

All selected primes are coprime to Nr for large N. They exceed $e^{L^\kappa} > r$, and their upper bound Q is smaller than $e^{\eta L}$. Consequently they are also coprime to u. Suppose that u is Q-smooth and that the residual scales satisfy

$$e^{3L/4} \le \frac{Z}{up} \le e^{5L/4} \qquad (p \in \mathcal{P}_0). \tag{6.9}$$

The common residual coefficient is

$$R_u(v) = \vartheta^{\omega_{\mathcal{P}_0}(v)} O(uv), \qquad |R_u(v)| \le L^{C_0}.$$

It depends on the Q-smooth part of v.

For a Dirichlet character χ modulo q, define

$$S_{u,j}(\chi) = \frac{1}{V} \sum_{p \in \mathcal{P}_j} \sum_{v \in \mathcal{V}_j} F(v)R_u(v)\, \frac{\psi(upv/Z)}{upv}\, e(\beta upv)(upv)^{i\zeta}\chi(p)\chi(v). \tag{6.10}$$

The denominator V in every slice is the mass of the complete group $\mathcal{P}_0$.

6.3. **Exact additive and multiplicative energy.** Define the corresponding additive polynomial on $\mathbb{Z}/qg\mathbb{Z}$ by

$$A_u(h) = \frac{1}{V} \sum_{j<K} \sum_{p \in \mathcal{P}_j} \sum_{v \in \mathcal{V}_j} F(v)R_u(v)W_{\beta,\zeta}(upv)e\Big(\frac{\overline{vpv}h}{q}\Big).$$

The factors p and v are units modulo q. Collecting all triples (j, p, v) with the same product residue gives a function c on $(\mathbb{Z}/q\mathbb{Z})^\times$. Repeated products contribute to the same value of

c, with their coefficients added. Additive orthogonality, multiplication by the unit v, and character orthogonality on the unit group give, successively,

$$\begin{aligned}\sum_{h \bmod qg} |A_u(h)|^2 &= gq \sum_{x \in (\mathbb{Z}/q\mathbb{Z})^\times} |c(x)|^2 \\ &= \frac{gq}{\varphi(q)} \sum_{\chi \bmod q} \left| \sum_{j<K} S_{u,j}(\chi) \right|^2 \\ &= \frac{gq}{\varphi(q)} \left\{ \left| \sum_{j<K} S_{u,j}(\chi_0) \right|^2 + \sum_{\chi \neq \chi_0} \left| \sum_{j<K} S_{u,j}(\chi) \right|^2 \right\}. \end{aligned} \tag{6.11}$$

This identity includes all cross terms between slices and between repeated product residues. For $q = 1$, its nonprincipal sum is empty.

6.4. **Mellin separation and the nonprincipal contribution.** We use the strong zero-free theorem of OpenAI, [7, Theorem 1.1], that the Dirichlet L-function of every character has no zero in $\operatorname{Re} s > 7/8$, with the principal pole at $s = 1$ excluded. The modulus is unrestricted. The logarithmic derivative and a smoothed Perron integral in this fixed half-plane give the following uniform prime estimate. For $q \le e^{2L}$, $|\tau| \le e^{3L}$, and $a_j \ge e^{L^\kappa}$,

$$\left| \sum_{p \in \mathcal{P}_j} \chi(p) p^{i\tau} \right| \ll a_j \exp(-L^\kappa/512) \qquad (\chi \neq \chi_0). \tag{6.12}$$

One way to obtain this form is first to prove

$$\left| \sum_{a < p \le b} \frac{\chi(p) p^{i\tau}}{p} \right| \ll (1 + \log((q+2)(|\tau|+3))) a^{-1/256},$$

and then to apply partial summation on $a_j < p \le b_j \le 2a_j$. The constants are uniform in the character, the modulus, and the twist in these ranges.

Put $\widetilde{\psi}(x) = \psi(x)/x$ and

$$\Phi_w(x) = \widetilde{\psi}(e^x) e(we^x), \qquad \widehat{\Phi}_w(t) = \int_{\mathbb{R}} \Phi_w(x) e(-tx)\, dx.$$

Fourier inversion in the logarithmic coordinate gives the exact formula

$$\begin{aligned} S_{u,j}(\chi) = \frac{1}{Z} \int_{\mathbb{R}} \widehat{\Phi}_{\beta Z}(t) Z^{-2\pi i t} u^{i(\zeta + 2\pi t)} &\left(\frac{1}{V} \sum_{p \in \mathcal{P}_j} p^{i(\zeta+2\pi t)} \chi(p) \right) \\ &\cdot \left(\sum_{v \in \mathcal{V}_j} F(v) R_u(v) v^{i(\zeta+2\pi t)} \chi(v) \right) dt. \end{aligned} \tag{6.13}$$

The outside factor $1/Z$ follows from $\psi(upv/Z)/(upv) = Z^{-1}\widetilde{\psi}(upv/Z)$. The factors involving only u and Z in the integrand have absolute value one. All sums are finite and $\widehat{\Phi}_{\beta Z}$ is integrable, so the interchanges in (6.13) are justified directly.

For coefficients b_v of absolute value at most L^{C_0} on an interval of length ℓ, unit-group orthogonality gives

$$\sum_{\chi \bmod q} \left| \sum_{v} b_v \chi(v) \right|^2 \le \varphi(q)(1 + \ell/q) \ell L^{2C_0}. \tag{6.14}$$

Indeed, each residue contains at most $1 + \ell/q$ interval points, and Cauchy–Schwarz applies separately in each residue. Using (6.7) and the height condition gives

$$(q + \ell_j)\ell_j \le 2\left(\frac{8Z}{ua_j}\right)^2.$$

Together with (6.12), this bounds the nonprincipal energy of the integrand in (6.13), after multiplication by $gq/\varphi(q)$, by

$$\frac{g}{u^2}\, L^{2C_0} V^{-2}\, O(\exp(-L^\kappa/256)).$$

The factors a_j^2 and Z^2 cancel against the interval length and the outside reciprocal normalization, respectively. Applying Cauchy–Schwarz to the sum of K slices costs at most K^2.

Enlarge the fixed exponent bounds to take $K_0, B_0 \ge 0$, with $K \le L^{K_0}$, $|\beta|Z \le L^{B_0}$ and $|\zeta| \le L^{B_0}$. Choose a fixed $D_0 > B_0 + 1$ and split the integral at $|t| = L^{D_0}$. On the central interval, $|\zeta + 2\pi t| \le e^{3L}$ for sufficiently large L. For every integer $k \ge 2$, compact support and differentiation give

$$\int_{|t|>T} |\widehat{\Phi}_w(t)|\, dt \ll_k (1 + |w|)^k T^{1-k}, \qquad \int_{\mathbb{R}} |\widehat{\Phi}_w(t)|\, dt \ll (1 + |w|)^2.$$

Thus the tail has any prescribed negative power of L after k is chosen sufficiently large. In the tail estimate use $|\sum_{p\in\mathcal{P}_j} \chi(p)p^{i\tau}| \le 2a_j$ and (6.14). The same cancellation of the physical length factors applies. Minkowski's inequality in the finite character space, followed by the displayed Fourier bounds, proves

$$\frac{gq}{\varphi(q)} \sum_{\chi\neq\chi_0} \left|\sum_{j<K} S_{u,j}(\chi)\right|^2 \ll_A \frac{g}{u^2} L^{-A} \qquad (A > 0). \tag{6.15}$$

The choices of D_0, k, and the threshold for L precede all finite coefficient and residue data.

6.5. **The principal character and the total residue bound.** For $p \in \mathcal{P}_j$ put $Z_p = Z/(up)$. The principal summand in (6.10) can be written as

$$\frac{(up)^{i\zeta}}{Z} \sum_{\substack{v\ge 1\\(v,Nr)=1}} F(v)R_u(v)\widetilde{\psi}(v/Z_p)e(\beta upv)v^{i\zeta}.$$

The positive unit indices at which the cutoff is nonzero lie in $\mathcal{V}_j$, as proved above. The principal cancellation theorem of Section 4 applies with scale Z_p. Its test has bounded support and polynomially bounded amplitude and variation, uniformly in p, by (6.9) and the fixed bounds on βZ and ζ. Its smooth-part coefficient is R_u, whose bound is (6.2). Increasing the accuracy in that theorem yields, for every $A_0 > 0$, a bound $L^{-A_0}/(up)$ for the absolute value of the last display. Consequently

$$|S_{u,j}(\chi_0)| \ll_{A_0} \frac{L^{-A_0}}{u}\frac{1}{V}\sum_{p\in\mathcal{P}_j}\frac{1}{p} \le \frac{L^{-A_0}}{u}. \tag{6.16}$$

Since $(N, r) = 1$, the roughness of N gives

$$\frac{q}{\varphi(q)} \le \frac{N}{\varphi(N)}\,\frac{r/g}{\varphi(r/g)} \le 2r \le 2L^D$$

for sufficiently large N. Combining (6.16), $K \le L^{K_0}$, and (6.15) in the exact identity (6.11), with A_0 chosen after the desired accuracy, proves the local estimate

$$\boxed{\sum_{h \bmod Nr} |A_u(h)|^2 \ll_A \frac{(u,r)}{u^2} L^{-A}.} \tag{6.17}$$

All powers lost in the principal and nonprincipal estimates are fixed before L tends to infinity.

6.6. **Aggregation and the deleted parts.** Let $U > 0$ and suppose that the smoothness of u, (6.8), and (6.9) hold uniformly for every supported part $u \le U$ retained in the physical sum. Sum over these supported parts. By the triangle inequality in $\ell^2(\mathbb{Z}/Nr\mathbb{Z})$,

$$\left(\sum_{h \bmod Nr} \left| \sum_u A_u(h) \right|^2 \right)^{1/2} \ll_A L^{-A/2} \sum_{\substack{u \ge 1 \\ p|u \Rightarrow p|r}} \frac{\sqrt{(u,r)}}{u}. \tag{6.18}$$

The last sum has a convergent Euler product. If $p^a \parallel r$, its local factor is

$$\sum_{k=0}^{a} p^{-k/2} + p^{a/2} \sum_{k>a} p^{-k} \le p^{2a}.$$

For $p \ge 3$ this follows by extending the first geometric sum to infinity and comparing with p^2. For $p = 2, a = 1$ the factor is $1 + \sqrt{2} < 4$, and for $a \ge 2$ it is bounded by $1/(1 - 2^{-1/2}) < 16 \le 2^{2a}$. Thus the product is at most r^2, and

$$\sum_{h \bmod Nr} \left| \sum_u A_u(h) \right|^2 \ll_A r^4 L^{-A}. \tag{6.19}$$

The discarded terms are controlled in physical coefficient space. Fix $M \ge 0$ and $C \in \mathbb{R}$, independently of L. For a coefficient sequence c_n supported on $1 \le n \le H_0$ with $|c_n| \le ML^C/Z$, Parseval gives

$$\int_{\mathbb{R}/\mathbb{Z}} \left| \sum_{\substack{n \le H_0 \\ (n,N)>1}} c_n e(n\theta) \right|^2 d\theta \le \frac{M^2 L^{2C}}{Z^2} H_0 \sum_{p|N} \frac{1}{p} \ll_A Z^{-1} L^{-A}, \tag{6.20}$$

$$\int_{\mathbb{R}/\mathbb{Z}} \left| \sum_{\substack{n \le H_0 \\ u_r(n) > U}} c_n e(n\theta) \right|^2 d\theta \le \frac{M^2 L^{2C}}{Z^2} H_0 U^{-1/2} r^2. \tag{6.21}$$

For the first estimate use $\sum_{p|N} p^{-1} \le \omega(N) e^{-\eta L}$. For the second, Rankin's inequality and the same supported-factor decomposition give

$$\#\{n \le H_0 : u_r(n) > U\} \le H_0 \sum_{\substack{u > U \\ p|u \Rightarrow p|r}} \frac{1}{u} \le H_0 U^{-1/2} \prod_{p|r} (1 - p^{-1/2})^{-1} \le H_0 U^{-1/2} r^2.$$

Taking $U = \exp(L^c/4)$ with fixed $c > 0$ makes (6.21) at most $Z^{-1}L^{-A}$ for every fixed A, after its threshold is increased. These estimates are uniform in the phases of the original coefficients.

6.7. **Returning to the original selected mark.** The selected first-order weight is

$$w_0(n) = \frac{\omega_{\mathcal{P}_0}(n)\vartheta^{\omega_{\mathcal{P}_0}(n)-1}}{V},$$

with the value zero when $\omega_{\mathcal{P}_0}(n) = 0$. The original endpoint coefficient is

$$c_n = F(n)w_0(n)O(n)W_{0,\zeta}(n).$$

Expanding the mark, and then using the invariance of F, gives

$$c_n = \frac{W_{0,\zeta}(n)}{V} \sum_{\substack{p\in\mathcal{P}_0\\ p\mid n}} F(n/p)\vartheta^{\omega_{\mathcal{P}_0}(n)-1}O(n).$$

If $p^2 \nmid n$, disjointness of the other prime groups gives $O(n) = O(n/p)$, and $\omega_{\mathcal{P}_0}(n) - 1 = \omega_{\mathcal{P}_0}(n/p)$. Replacing the residual factor by $\vartheta^{\omega_{\mathcal{P}_0}(n/p)}O(n/p)$ therefore changes only the terms with $p^2 \mid n$. The resulting common coefficient is the one used in (6.10).

This replacement is valid with every retained mask inside the energy. Here is a direct quantitative bound. Write the selected-prime average with probability weights $1/(Vp)$, and consider its fixed-p error. Its absolute value is at most $4pL^{C_0}/Z$, it is supported on multiples of p^2, and $(p^2, Nr) = 1$. Among these multiples, each residue modulo Nr occurs at most $1+H_0/(p^2Nr)$ times. By (6.8), this is at most $2H_0/(p^2Nr)$. Cauchy–Schwarz in each residue and then additive orthogonality give a fixed-p residue energy at most

$$\frac{32H_0^2L^{2C_0}}{p^2Z^2} \le \frac{2048L^{2C_0}}{p^2}.$$

Jensen's inequality for the probability weights $1/(Vp)$ bounds the energy of the full replacement by

$$\frac{4096L^{2C_0}}{P_{\min}^2}, \qquad P_{\min} = \min \mathcal{P}_0 \ge e^{L^\kappa}. \tag{6.22}$$

The argument also applies after an additive modulation and after multiplication by any mask of absolute value at most one. Thus (6.22) has arbitrary logarithmic saving. The elementary inequality $|z + w|^2 \le 2|z|^2 + 2|w|^2$, together with (6.19) at a higher accuracy, transfers the retained residue estimate to the original marked endpoint.

Finally, $w_0(n) \le 4$ follows from $V \ge 1/2$ and $0 \le \vartheta \le 1/2$, so $|c_n| \le 8L^{C_0}/Z$. The deletion estimates (6.20) and (6.21) consequently apply to this original endpoint. Let $\mathfrak{M}_r$ be the union of the neighborhoods of radius Δ_r about the rationals with denominator Nr. For the retained coefficient, integrating the complete residue estimate over $|\beta| \le \Delta_r$ bounds its energy on $\mathfrak{M}_r$ by $4\Delta_r E_r$, where E_r is the uniform residue-energy bound. Write $c_{\mathrm{ret},r}$, c_{bad}, and $c_{\mathrm{large},r}$ for the three coefficients in the exact decomposition $c = c_{\mathrm{ret},r} + c_{\mathrm{bad}} + c_{\mathrm{large},r}$. Subadditivity over denominator sets and the three-term square inequality give

$$\int_{\bigcup_{r\le R}\mathfrak{M}_r} |\widehat{c}(\theta)|^2\, d\theta \le 12\sum_{r\le R}\Delta_r E_r + 3RE_{\mathrm{bad}} + 3\sum_{r\le R} E_{\mathrm{large},r}.$$

Here each deleted term is restricted to the whole set $\mathfrak{M}_r$ before its integral is bounded by its full-circle energy. This uses no multiplicity factor from the individual rational neighborhoods. With R and $Z\Delta_r$ bounded by fixed powers of L, increasing the local accuracies proves an upper bound $Z^{-1}L^{-A}$ for every prescribed $A > 0$.

7. The comparison terms and the centered correlation

Throughout this section, $L = \log N$, the integer N is N^{ε}-rough, and $0 < \varepsilon < 1/4$ is fixed. We use the groups and the signed family of Section 5. Write $\vartheta = 1/256$ for the damping parameter, $Q = \exp(L^{.49})$ for the smoothness threshold, and

$$Q_0 = \lceil L^{q_*} \rceil, \qquad h = L^{-h_*}, \qquad \xi = L^{-x_*}$$

for the parameters of the product-cell kernels. The function F_0 is the real invariant center of Section 4. In particular, $|F_0| \le 1$ and $F_0(dn) = F_0(n)$ for every positive integer n and every positive Q-smooth integer d.

All prime labels below have their original probability weights $\mu_g(p) = 1/(V_g p)$. Every sum over free labels includes repeated primes. The injectivity condition on an active endpoint is retained until the explicit restriction estimate below. We extend a function on positive integers to the negative endpoint by putting $\widetilde{H}(-n) = H(n)$ for $n > 0$ and $\widetilde{H}(n) = 0$ for $n \ge 0$.

7.1. **A common target and the reciprocal normalization.** Let $X = 2Y$ be an even positive integer and put

$$T_X = [-6X, -1] \cap \mathbb{Z}.$$

We take a real smooth function ψ, supported in $(1/2, 2)$, with $|\psi| \le 1$. The soft source is $F_0(x)\psi(x/X)$ for $x > 0$. Its two hard source intervals are $[Y, 2Y)$ and $[X, 2X)$, and both use the target T_X.

If $x > 0$, $\chi(x/(x-y)) \ne 0$, and $\operatorname{supp}\chi \subset [1/4, 3/4]$, then $y < 0$ and $|y| \le 3x$. Consequently $x < 2X$ implies $y \in T_X$. Thus the projection to T_X may be inserted into every cutoff pairing on this soft source. This projection is kept in the ensuing frequency integrals. For a hard source scale Y it has radius $12Y$, and for scale X it has radius $6X$.

For a single pattern ν, write $\mathcal{R}_{\nu,X}(\zeta)$ for its normalized physical pairing, with step $-N$, additive parameter zero, source

$$d_\vartheta(x)\frac{X}{x}F(x), \qquad F(x) = F_0(x)\psi(x/X)(x/N)^{i\zeta},$$

and target $d_\vartheta(y)\mathbf{1}_{T_X}(y)\widetilde{H}(y)$. The pairing is divided by X. Inner products are conjugate-linear in the first variable. The definitions also apply to a hard source restriction of F. For $\zeta = -2\pi t$ this is the frequency that occurs in the Mellin formula below.

Fix a comparison pattern $\nu = (A, B)$. In group g, put

$$s_g = |S_{\nu,g}|, \qquad c_g = |C_{\nu,g}|, \qquad t_g = 1 + c_g = M - s_g.$$

A shared tuple $\mathbf{s}$ has product D and probability

$$D = \prod_{g, i \in S_{\nu,g}} s_{g,i}, \qquad \mu_{\mathrm{sh}}(\mathbf{s}) = \prod_{g, i \in S_{\nu,g}} \frac{1}{V_g s_{g,i}} = \frac{\sigma_\nu}{D}, \qquad \sigma_\nu = \prod_g V_g^{-s_g}. \tag{7.1}$$

The empty tuple has product and probability one. The two free tuples α, β occupy A, B, have their product laws μ_A, μ_B, and have products $a = a(\alpha)$ and $b = b(\beta)$. These two tuples exhaust the replaced coordinates. Therefore the edge equations are exactly

$$x = Dn, \qquad y = D(n - Nab). \tag{7.2}$$

For an integer $n \ne 0$, let $\Lambda_\nu(n)$ be the finite set of ordered lists which choose, in each group g, t_g distinct primes dividing $|n|$. For a function E of those lists, define

$$\mathcal{G}_E(n) = \sum_{\lambda \in \Lambda_\nu(n)} E(\lambda) \prod_g \frac{\vartheta^{\omega_g(n) - t_g}}{V_g^{t_g}}. \tag{7.3}$$

An empty list sum is zero. Equivalently, one can sum over the fixed finite space of all ordered distinct prime lists and insert the condition that every label divides $|n|$. This equivalence will permit us to take the complete endpoint sum before a square norm.

Here is the normalization in the change of variables. Suppose the shared labels are distinct and are coprime to both residual endpoints. The number of group-g prime factors of Dn is $s_g + \omega_g(n)$, and the corresponding statement holds at $D(n - Nab)$. The two damping factors from the test vectors and the two from the edge give

$$\vartheta^{\sum_g (\omega_g(n) - t_g)} \vartheta^{\sum_g (\omega_g(n - Nab) - t_g)}.$$

The mass and the edge normalization satisfy

$$\prod_g V_g^{-M} V_g^{-t_g} = \sigma_\nu \prod_g V_g^{-2t_g}. \tag{7.4}$$

The reciprocal source then gives the exact identity

$$\frac{1}{X} \sigma_\nu \frac{X}{Dn} = \frac{\mu_{\mathrm{sh}}(\mathbf{s})}{n}. \tag{7.5}$$

In particular, the residual source carries $1/n$.

The physical endpoint lists decompose uniquely into their prescribed shared coordinates and their residual ordered lists. Conversely, these data reconstruct the two physical endpoint lists whenever the shared labels are distinct, are coprime to the residual endpoints, and the free labels avoid both endpoint lists. These are precisely the good-label conditions in this reindexing. The coefficient in (5.4) becomes

$$-(-1)^{|A|+|B|} \mathcal{K}_A\big(a, P_A(\lambda_{\mathrm{target}})\big) \mathcal{K}_B\big(b, P_B(\lambda_{\mathrm{source}})\big). \tag{7.6}$$

The first kernel is evaluated at the target list, and the second at the source list. The remaining phase is $(Dab)^{i\zeta}$.

These identities apply to the endpoint-permutation averages as well. Changing variables by a slot permutation preserves the physical vertex set, its mass, and both endpoint functions, which depend only on the integer position. Hence each permuted pairing equals its unpermuted counterpart.

7.2. **Removing the finite restrictions.** We record bounds that apply to a hard source $[X, 2X)$, then add the two hard pieces of a soft source. Set

$$\mathcal{E}_\nu = \prod_g \frac{t_g!}{(1 - \vartheta)^{t_g} V_g^{t_g}}, \qquad S_\nu = \sum_g s_g, \qquad C_\nu = \sum_g c_g.$$

The total absolute residual list weight at either endpoint is at most $\mathcal{E}_\nu$. Indeed,

$$(w)_t \vartheta^{w-t} = t! \binom{w}{t} \vartheta^{w-t} \le \frac{t!}{(1 - \vartheta)^t} \quad (w \ge t),$$

by the binomial theorem. Here $(w)_t = w(w-1)\cdots(w-t+1)$.

Let P be a common lower bound for the primes and $D_{\max}$ an upper bound for every shared product. Suppose the two kernel factors are bounded by M_A, M_B and the two position functions by B_F, B_H. Before division by X, the difference between the retained residual sum and the unrestricted residual sum has absolute value at most

$$\mathcal{E}_\nu^2 B_F M_A M_B B_H \left\{ (C_\nu + 2S_\nu) \left(\frac{2X}{P} + D_{\max} \right) + 2X \sum_g \frac{s_g^2 + s_g c_g}{P V_g} \right\}. \tag{7.7}$$

The unrestricted sum includes every shared tuple with the law (7.1), and every free tuple with its original law.

To see the terms in (7.7), fix the shared and free tuples. A forbidden incidence at either residual endpoint imposes one congruence modulo a prime $p \ge P$ on an interval of length at most $2X/D$. Its number of solutions is at most $2X/(Dp)+1$. The absolute summand, after the reciprocal is bounded on the hard source, is at most $DB_FM_AM_BB_H\mathcal{E}_\nu^2$. Thus one such incidence contributes at most the common factor times $(2X/P+D)$. There are at most $C_\nu+2S_\nu$ incidences of this kind. A coincidence between prescribed labels has probability at most

$$\sum_{p\in\mathcal{P}_g}\mu_g(p)^2 \le \frac{1}{PV_g}.$$

Summing over the possible shared–shared and shared–free pairs gives the last term of (7.7). A union bound is valid when several of the incidences occur simultaneously.

There is also the part of the physical pairing in which a shared prime divides a residual endpoint. Its normalized absolute value is

$$\ll M_AM_BB_FB_H\,T_\nu\sqrt{\Delta_\nu}, \qquad \Delta_\nu=\sum_g\frac{s_g}{V_g}\sum_{p\in\mathcal{P}_g}\frac{1}{p^2}\le\frac{S_\nu}{P}, \tag{7.8}$$

where T_ν is the physical row bound for this pattern. For completeness, the squared mass of the source restricted to these incidences is at most a fixed multiple of $B_F^2X\Delta_\nu$. This follows by counting the multiples of the shared label once in the shared harmonic average and once in the position sum. The same assertion holds for the target. Apply Cauchy–Schwarz to the source-incidence and target-incidence parts separately, use the row bound for the edge, and add them. The unrestricted source and target squared masses are bounded by fixed multiples of X, which proves (7.8).

For the actual family, $t_g\le m+1$, $V_g\ge 1/2$ and $|G|\le 2\log L$. Thus $\mathcal{E}_\nu$ and T_ν are bounded by fixed powers of L. Also $M_A,M_B\le Q_0^2/\xi$, $S_\nu\le 2J\log L$, $C_\nu\le m|\mathcal{B}|$, and, for every fixed J,

$$P\ge\exp(L^{.27}), \qquad 1\le D\le D_{\max}\le\exp(L^{.60}) \tag{7.9}$$

eventually. Divide (7.7) by X, and use $X\ge\exp(.75L)$. Together with (7.8), this gives

$$|\mathcal{R}_{\nu,X}-\mathcal{U}_{\nu,X}|\le C_JL^{C_*}\left\{\exp(-\tfrac12L^{.27})+\exp(L^{.60}-.75L)\right\}. \tag{7.10}$$

Here $\mathcal{U}_{\nu,X}$ is the normalized unrestricted expression, C_* depends only on the fixed construction parameters, and C_J may depend on J. In particular the error is $O(L^{-A})$ for every fixed A. For the raw branch the same reindex uses empty selected sets and auxiliary kernel factors -1 and 1. Their product with $-(-1)^0$ is the raw coefficient one. Thus the same conclusion holds for that branch. Both estimates hold uniformly in ζ.

7.3. **Finite features and complete endpoint sums.** For a nonempty selected set C, let $\mathscr{F}_C$ consist of the pairs (j,χ) such that j is an occupied original product cell and $\chi\in\mathcal{C}(Q_0)$. Define

$$\begin{aligned} u_{C,j,\chi}(v)&=\mathbf{1}_{j(v)=j}\overline{\chi(v)},\\ v_{C,j,\chi}(v)&=\mathbf{1}_{j(v)=j}\mathbf{1}_{\nu_{C,j}\ge\xi}\nu_{C,j}^{-1}\chi(v). \end{aligned} \tag{7.11}$$

Terms with $\nu_{C,j}=0$ are zero. Then

$$\mathcal{K}_C(d,v)=\sum_{\lambda\in\mathscr{F}_C}u_{C,\lambda}(d)v_{C,\lambda}(v), \qquad |u_{C,\lambda}|\le1, \quad |v_{C,\lambda}|\le\xi^{-1}. \tag{7.12}$$

This is a finite identity obtained directly from (5.5).

For $\lambda\in\mathscr{F}_A$ and $\rho\in\mathscr{F}_B$, form the complete endpoints

$$G_{A,\lambda}(n)=\mathcal{G}_{v_{A,\lambda}\circ P_A}(n), \qquad G_{B,\rho}(n)=\mathcal{G}_{v_{B,\rho}\circ P_B}(n).$$

Fix a shared tuple, put $Z = X/D$, and set

$$f_{\rho,D}(n) = \frac{\overline{F(Dn)}}{n} G_{B,\rho}(n),$$
$$g_{\lambda,D}(r) = G_{A,\lambda}(r)\mathbf{1}_{T_X}(Dr)\widetilde{H}(Dr). \tag{7.13}$$

The source is supported in $(Z/2, 2Z)$, and the target in $[-6Z, 0)$. For a pair of features define

$$\mathcal{C}_{D,\lambda,\rho} = \sum_{\alpha,\beta} \mu_A(\alpha)\mu_B(\beta)a^{i\zeta}u_{A,\lambda}(a)b^{i\zeta}u_{B,\rho}(b) \sum_n f_{\rho,D}(n)g_{\lambda,D}(n - Nab). \tag{7.14}$$

The coefficient dictionary, the two finite feature expansions, and (7.5) give

$$\mathcal{U}_{\nu,X}(\zeta) = -(-1)^{|A|+|B|} \sum_{\mathbf{s}} \mu_{\mathrm{sh}}(\mathbf{s})D^{i\zeta} \sum_{\lambda\in\mathscr{F}_A,\, \rho\in\mathscr{F}_B} \mathcal{C}_{D,\lambda,\rho}. \tag{7.15}$$

All rearrangements here are finite. In particular, $G_{A,\lambda}$ and $G_{B,\rho}$ include every residual label before any analytic estimate is applied.

We will use their full square masses. Put

$$C_m = (m+1)!4^{m+1}, \qquad B_2 = 2x_* + 2\log C_m + 1.$$

For any endpoint feature E with $|E| \le \xi^{-1}$, let $W_\nu(n)$ be the endpoint sum (7.3) with $E = 1$. The binomial bound gives $W_\nu(n) \le C_m^{|G|}$ and $|\mathcal{G}_E(n)| \le \xi^{-1}W_\nu(n)$. Moreover, for $T \ge 1$,

$$\sum_{0<|n|\le T} W_\nu(n) \le 2T. \tag{7.16}$$

Indeed, after fixing a distinct list with product d, its multiples contribute at most $2T/d$, and its damping is at most one. Summing $1/d$ over the lists gives at most $\prod_g V_g^{t_g}$, cancelling the normalizing denominator. Consequently

$$\sum_{0<|n|\le T} |\mathcal{G}_E(n)|^2 \le 2T\xi^{-2}C_m^{|G|}.$$

Since $1/n \le 2/Z$ on the source, this proves, for sufficiently large L,

$$\sum_n |f_{\rho,D}(n)|^2 \le Z^{-1}L^{B_2}, \qquad \sum_r |g_{\lambda,D}(r)|^2 \le ZL^{B_2}. \tag{7.17}$$

For example, bounding the positive source height by $3Z$ gives $24\xi^{-2}C_m^{|G|}/Z$, and bounding the target height by $7Z$ gives $14\xi^{-2}C_m^{|G|}Z$. The extra power in B_2 absorbs these constants.

Choose one group $\mathcal{P}_0$ in the low block. It is untouched by every comparison pattern, so $t_0 = 1$, and neither endpoint feature depends on its label. The complete source factors as

$$G_{B,\rho}(n) = w_0(n)O_\rho(n), \qquad w_0(n) = \frac{\omega_{\mathcal{P}_0}(n)\vartheta^{\omega_{\mathcal{P}_0}(n)-1}}{V_0}. \tag{7.18}$$

The coefficient O_ρ keeps all the remaining group sums together. It depends only on the Q-smooth part of n and is bounded by a fixed power of L. By the invariance of F_0, the source in (7.13) is, up to a constant of modulus one,

$$F_0(n)w_0(n)O_\rho(n)\frac{\psi(n/Z)}{n}n^{i\tau}, \qquad |\tau| \le L^2. \tag{7.19}$$

The selected-prime expansion in Section 6 therefore has dilation one. All other marked factors remain in O_ρ.

7.4. **The product dyads and their major and minor arcs.** Partition the free product values into intervals $[U, 2U)$ and $[V, 2V)$, where U, V are powers of two. Every summand of (7.14) in such a pair has $n \in (Z/2, 2Z)$ and $n - Nab \in [-6Z, 0)$. Hence a nonzero product dyad satisfies

$$\frac{Z}{8} < NUV < 8Z. \tag{7.20}$$

For example, $Nab = n - (n - Nab)$ lies between $Z/2$ and $8Z$, while $UV \le ab < 4UV$. Every dyad failing (7.20) contributes zero.

Both A and B are nonempty and belong to the high groups. Thus each sampled free product is at least $\exp(L^{.39})$. For a fixed $c_f < .39$, say $c_f = 3/8$, eventually

$$U, V \ge \exp(L^{c_f}), \qquad ab \le \exp(L^{.60}), \qquad 1 + \log(2UV) \le L^{d_f}, \quad d_f = .48. \tag{7.21}$$

The upper bounds follow from at most m selected slots in each high group and $\log p \le L^{.46}$. They are uniform over the selected sets. The same estimates for the shared product put $Z = X/D$ in $[\exp(.75L), \exp(1.25L)]$ whenever the physical soft scale is in $[\exp(.9L), \exp(1.1L)]$.

There is a fixed $C_a \ge 0$, depending on m, such that the original product distributions satisfy

$$\sum_{\alpha:\, a(\alpha)=a} \mu_A(\alpha) \le \frac{L^{C_a}}{a}, \qquad \sum_{\beta:\, b(\beta)=b} \mu_B(\beta) \le \frac{L^{C_a}}{b}. \tag{7.22}$$

Unique factorization determines the multiset in each group from the product. At most $m!$ assignments yield this multiset, and their mass is at most 2^m times the reciprocal of its product. Since $|\mathcal{B}| \le .07 \log L$, one can take any $C_a > .07 \log(m! 2^m)$. This argument includes products with repeated prime factors. Restriction to a product dyad gives atom bounds L^{C_a}/U and L^{C_a}/V, and total masses at most one.

Fix a desired minor-arc exponent $B > 0$ and put

$$\Gamma = 2C_a + 2B + 8, \qquad R_0 = \lceil L^{\Gamma} \rceil, \qquad S = L^{d_f + \Gamma}, \qquad \Delta_r = \frac{S}{NrUV} \quad (1 \le r \le R_0). \tag{7.23}$$

Let $\mathfrak{M}$ be the union, on $\mathbb{R}/\mathbb{Z}$, of the arcs of radius Δ_r around every rational with denominator Nr, for $1 \le r \le R_0$. The free-label multiplier of the dyad is

$$M(\theta) = \sum_{\alpha,\beta} \mu_A(\alpha)\mu_B(\beta) a^{i\zeta} u_{A,\lambda}(a) b^{i\zeta} u_{B,\rho}(b) e(-Nab\theta),$$

where the two sums are restricted to their dyads. Its absolute value is at most one everywhere.

Lemma 7.1. *For sufficiently large L, $|M(\theta)| \le L^{-B}$ on $\mathfrak{M}^c$, uniformly in both endpoint features and ζ.*

Proof. Put $T = UV/S$, which tends to infinity. Dirichlet approximation gives coprime integers a_0, r, with $1 \le r \le T$, such that

$$|N\theta - a_0/r| \le \frac{1}{rT}, \qquad |N\theta - a_0/r| \le \frac{1}{r^2}.$$

If $r \le R_0$, division by N puts θ in the arc of radius $S/(NrUV)$ about $a_0/(Nr)$. Hence $r > R_0$ on the minor arcs.

We recall the elementary weighted bilinear bound in this normalization. For product distributions of mass at most one, supported in $[U, 2U)$ and $[V, 2V)$, with atoms at most K/U and K/V, and for unit coefficients, a reduced approximation $|\alpha - a_0/r| \le r^{-2}$ gives

$$\left| \sum_{a,b} c_a d_b e(\alpha ab) \right| \le 4K \left\{ \frac{1}{V} + \frac{1}{r} + \frac{1 + \log(2r)}{U} + \frac{r(1 + \log(2r))}{UV} \right\}^{1/2}. \tag{7.24}$$

One obtains this by Cauchy–Schwarz in the a variable. Expanding the square in b, the geometric sum in a is bounded by $\min(2U, \|\alpha(b-b')\|^{-1})$. Sorting the differences modulo r and summing the reciprocal distances j/r gives the last three terms in braces. The diagonal gives $1/V$. The atom bounds supply the two factors of K before the square root. The same proof first groups the original tuples by product, so it applies to our distributions.

Apply (7.24) with $\alpha = -N\theta$ and $K = L^{C_a}$. The two terms involving $1/U$ or $1/V$ are smaller than every fixed negative power of L. The other terms are bounded by $L^{-\Gamma}$, since $r > R_0$, $r \le UV/S$, and $1 + \log(2r) \le L^{d_f}$. The margin eight in Γ absorbs the constant and proves the claim. □

Lemma 7.2. *For every fixed $A_1 > 0$, the complete source in (7.19) satisfies*

$$\int_{\mathfrak{M}} |\widehat{f}_{\rho,D}(\theta)|^2 \, d\theta \le Z^{-1} L^{-A_1} \tag{7.25}$$

for sufficiently large L, uniformly in the shared tuple, the two features, the retained product dyad, and $|\zeta| \le L^2$.

Proof. We verify the parameters for the complete selected-mark estimate in Section 6. The chosen low group has

$$V_0 \ge \tfrac{1}{2}, \qquad \exp(L^{.27}) \le p \le \exp(L^{.36}) \quad (p \in \mathcal{P}_0).$$

All its primes are coprime to Nr, since N is N^{ε}-rough and $r \le R_0$ is bounded by a fixed power of L. Keep the canonical r-supported part u only when $u \le U_0 = \exp(L^{.39}/4)$. With $H_0 = \lfloor 8Z \rfloor$, we have

$$p^2 u N r \le H_0. \tag{7.26}$$

Indeed $NUV < 8Z$ by (7.20), whereas $p^2 ur/(UV)$ tends to zero uniformly. The sharper lower bound $UV \ge \frac{1}{4}\exp(2L^{.39})$ bounds its logarithm by $2L^{.36} - \frac{7}{4}L^{.39} + O(\log L)$, which tends to minus infinity. Thus the left side of (7.26) is eventually at most $4Z$, and its integrality gives the displayed inequality. Likewise $Z/(up) \in [\exp(.75L), \exp(1.25L)]$ eventually, after using the interior physical scale and (7.9).

Delete the N-nonunits and the terms whose r-supported part exceeds U_0. Their full-circle energies are $Z^{-1}L^{-A'}$ for every fixed A', by (6.20) and (6.21). The coefficient cap required there follows from (7.18). These deletions occur in the complete source coefficient.

For a retained u, write the positive source integer as upv, with $(v, Nr) = 1$, and set $g = (u, r)$ and $q_r = N(r/g)$. Since $(N, r) = 1$, the integer u/g is a unit modulo q_r, and

$$e\left(\frac{upva_0}{Nr}\right) = e\left(\frac{(u/g)pva_0}{q_r}\right).$$

Every residue modulo q_r occurs g times as a_0 runs modulo Nr. The exact character identity (6.11) therefore has the factor $gq_r/\varphi(q_r)$. The principal character is bounded by Theorem 4.9 and Corollary 4.10, with dilation one. The remaining coefficient is the entire O_ρ. For a nonprincipal character, use the dyadic prime slices and the reciprocal Mellin identity (6.13). Its external factor is $1/Z$. The original V_0 normalizes every slice. Corollary 3.6 applies because $q_r \le e^{2L}$ and the truncated frequencies are at most e^{3L}. The fixed smooth profile has arbitrarily rapid Fourier-tail decay. The residual orthogonality estimate (6.14) controls that tail with the same quotient factor.

It follows from (6.17) that the complete residue energy for this u is at most $(u, r)u^{-2}L^{-A'}$ for any prescribed A'. Minkowski and (6.18) give the full retained energy at most $r^4 L^{-A'}$. The selected-prime square replacement is valid with both deletion masks inside its coefficient, by (6.22) and (7.26).

Finally, (7.20) and (7.23) give $Z\Delta_r \le 8S/r$. Thus both R_0 and $Z\Delta_r$ have fixed polynomial bounds. The major-set cover and the three-part deletion inequality at the end of Section 6, with A' increased, prove (7.25). This accounts for the full residue multiplicity, the totient factor, and every prime slice. □

For finite sequences use the Fourier convention $\widehat{f}(\theta) = \sum_n f(n)e(n\theta)$. Finite orthogonality gives the exact identity

$$\mathcal{C}^{U,V}_{D,\lambda,\rho} = \int_{\mathbb{R}/\mathbb{Z}} \widehat{f}_{\rho,D}(\theta)\widehat{g}_{\lambda,D}(-\theta)M(\theta)\,d\theta. \tag{7.27}$$

The superscript denotes restriction to the two free product dyads. Apply Cauchy–Schwarz on $\mathfrak{M}$ and $\mathfrak{M}^c$, using $|M| \le 1$ on the former and Lemma 7.1 on the latter. Parseval and (7.17) yield

$$|\mathcal{C}^{U,V}_{D,\lambda,\rho}| \le L^{(B_2-A_1)/2} + L^{B_2-B}. \tag{7.28}$$

Both exponents can be made arbitrarily negative after the kernel parameters have been fixed.

7.5. **Summing every feature and every comparison pattern.** For a selected product, $\log a \le C'_m \log L\, L^{.46}$. It follows that the number of occupied product dyads is at most $L^{1.46}$. The number of occupied logarithmic cells is at most $C'_m \log L\, L^{.46}/h + 2$. Since $|\mathcal{C}(Q_0)| \le Q_0^2$, eventually

$$|\mathscr{F}_C| \le L^{h_*+2q_*+.46+4}. \tag{7.29}$$

These counts are for the full product spaces. The mass threshold ξ only removes terms from the endpoint features.

Put

$$C_{\mathrm{sum}} = 2m\log 2 + 2(h_* + 2q_* + .46 + 4) + 2(1.46) + 1.$$

The number of nonraw patterns is at most $2^{m|\mathcal{B}|} \le L^{2m\log 2}$. Thus the product of this count, both feature counts, and both dyad counts is at most $L^{C_{\mathrm{sum}}-1}$. Given $A > 0$, choose $B \ge B_2 + A + C_{\mathrm{sum}} + 3$ and $A_1 \ge B_2 + 2(A + C_{\mathrm{sum}} + 3)$ in (7.28). Sum it over all features and all retained product dyads. Then average over the shared law, whose mass is exactly one. The phase $D^{i\zeta}$ and the sign in (7.15) have modulus one. This proves

$$\sum_{\nu\neq 0} |\mathcal{U}_{\nu,X}(\zeta)| \le L^{-A-1} \tag{7.30}$$

eventually. Adding (7.10) for the two hard source pieces and all patterns gives

$$\sum_{\nu\neq 0} |\mathcal{R}_{\nu,X}(\zeta)| \le L^{-A} \qquad (|\zeta| \le L^2). \tag{7.31}$$

All choices in this summation occur after m, q_*, h_*, x_* and J. They change the arithmetic precision and its threshold. The graph family and its frequency range stay fixed.

7.6. **Using the full physical moment.** We spell out the moment-to-pairing step with the common target. Let a finite weighted graph have integer positions, and suppose every edge has displacement kd with $k \neq 0$ and $|d| \le \mathcal{B}$. For its position roots u_n, let $T_R = \sum_{n\in I}\langle u_n, (AA^*)^R u_n\rangle$. If $f = \sum_{n\in I} a_n u_n$ and $|a_n| \le 1$, then

$$|\langle f, Ag\rangle| \le \sqrt{3}\,(\mathcal{B}+2)^{1/(2R)}\|g\|\,\|f\|^{1-1/R}T_R^{1/(2R)}. \tag{7.32}$$

To prove this, partition each residue class modulo $|k|$ into successive blocks of $b = \lceil \mathcal{B} \rceil$ positions. Write f_j for the source in a block and restrict g to that block and its two neighbors.

Every coordinate of g belongs to at most three such neighborhoods. Spectral Hölder for the positive operator AA^* gives

$$\|A^*f_j\| \le \|f_j\|^{1-1/R}\langle f_j, (AA^*)^R f_j\rangle^{1/(2R)}.$$

Cauchy–Schwarz in the at most b roots of a block bounds its last quadratic form by b times the sum of its individual root moments. Hölder in the blocks, with exponents 2, $2R/(R-1)$ and $2R$, now proves (7.32). Orthogonality of the position fibers gives $\sum_j \|f_j\|^2 = \|f\|^2$, and the three-neighborhood overlap gives the factor $\sqrt{3}$.

Apply this to a hard source $[Z_0, 2Z_0)$ of the actual physical graph. The root coefficient is $d_\vartheta(n)(Z_0/n)F(n)$, of modulus at most one on the physical vertices. Counting the ordered endpoint labels by their products gives source and common-target squared masses at most $24Z_0$. Thus both norms are at most $5\sqrt{Z_0}$. The constructed moment theorem gives $T_R \le Z_0(L^{-E})^{2R}$. Substitution into (7.32), followed by division by Z_0, yields

$$|\mathcal{R}_{\mathrm{full},Z_0}(\zeta)| \le 25\sqrt{3}\,(\mathcal{B}+2)^{1/(2R)}L^{-E}. \tag{7.33}$$

Here $R = \lceil L^{.52}\rceil$ and $\log\mathcal{B} = O_J(L^{.46}\log L)$, so the middle factor is at most two eventually. The powers of Z_0 cancel exactly.

For the soft source, add the estimates at $Z_0 = Y$ and $Z_0 = X$, using T_X in both. The identities $Y^{-1}(Y/x) = X^{-1}(X/x) = 1/x$ show that their sum is the original normalized soft pairing. In particular,

$$|\mathcal{R}_{\mathrm{full},X}(\zeta)| \le 100\sqrt{3}L^{-E} \qquad (|\zeta| \le L^2). \tag{7.34}$$

The source cutoff in the arithmetic argument is still $\psi(x/X)$. The two hard restrictions have been summed before (7.19) is used.

Choose $E = 30$ before invoking Theorem 5.6. That theorem chooses the kernel parameters and then a lower bound for J. Fix J afterward. Equation (7.34) is then at most L^{-27} eventually. Put $C = 2m\log 2$, and choose the arithmetic precision so that every comparison term is at most $L^{-(27+C)}$. There are at most L^C such terms, so their sum is at most L^{-27}. Choose each of the two hard raw restriction errors at precision L^{-27} as well. The full graph is the sum of its raw and comparison branches. Hence the normalized unrestricted raw expression satisfies

$$|\mathcal{U}_{0,X}(-2\pi t)| \le 4L^{-27} \le L^{-26} \qquad (|t| \le L) \tag{7.35}$$

for sufficiently large L. The fixed choice $|t| \le L$ is valid because $2\pi L \le L^2$ eventually.

7.7. **Mellin inversion and the global dyadic sum.** Put

$$W(n) = \prod_g \frac{\omega_g(n)\vartheta^{\omega_g(n)-1}}{V_g} \quad (n \ge 1), \tag{7.36}$$

where a factor with $\omega_g(n) = 0$ is zero. The raw branch has no free labels and has $t_g = 1$ in every group. Its unrestricted expression therefore simplifies to

$$\mathcal{U}_{0,X}(-2\pi t) = \sum_{\mathbf{s}} \mu_{\mathrm{sh}}(\mathbf{s}) \sum_{n\ge 1} \frac{W(n)W(|n-N|)}{n} F_0(n)\psi(Dn/X)\mathbf{1}_{T_X}(D(n-N))$$
$$\cdot\, \widetilde{H}(D(n-N))e\big(t\log(n/N)\big). \tag{7.37}$$

The value at $n = N$ is zero. In obtaining this identity, the shared factor $D^{-2\pi it}$ multiplies the conjugated source phase $(Dn/N)^{2\pi it}$, giving $(n/N)^{2\pi it}$.

Let χ be a fixed real smooth function supported in $(1/4, 3/4)$, and put

$$\Phi(u) = e^u\chi(e^u), \qquad \widehat{\Phi}(t) = \int_{\mathbb{R}} \Phi(u)e(-tu)\,du. \tag{7.38}$$

This is a smooth compactly supported function on $\mathbb{R}$. Fourier inversion gives

$$\int_{\mathbb{R}} \widehat{\Phi}(t) e\big(t \log(n/N)\big)\, dt = \frac{n}{N}\chi(n/N).$$

Multiplying (7.37) by $\widehat{\Phi}(t)$ and integrating is legitimate because its sums are finite. On the support of the resulting expression, $N/4 < n < 3N/4$, and $\psi(Dn/X) \neq 0$ implies $Dn < 2X$. Hence $-6X < D(n-N) < 0$, so the common-target indicator is one. If H is Q-smooth-invariant, then $H(D(N-n)) = H(N-n)$ as well. We obtain

$$\begin{aligned}\int_{\mathbb{R}} \widehat{\Phi}(t)\mathcal{U}_{0,X}(-2\pi t)\, dt = \frac{1}{N}\sum_{0<n<N} \chi(n/N)W(n)W(N-n)F_0(n)H(N-n) \\ \cdot \sum_{\mathbf{s}} \mu_{\mathrm{sh}}(\mathbf{s})\psi(Dn/X).\end{aligned} \tag{7.39}$$

This also proves directly why the positive-ratio profile is $v\chi(v)$.

We bound the complementary frequencies using the same raw expression. Let $M_1 = 4^{|G|}$, so that $W(n) \le M_1$. For fixed D, the nonzero source integers lie in $(Z/2, 2Z)$, where $1/n \le 2/Z$. There are at most $2Z$ such positive integers when $Z \ge 1$. The shared law has total mass one, and consequently

$$|\mathcal{U}_{0,X}(-2\pi t)| \le 4M_1^2 \le L^{4\log 4+2} \quad (t \in \mathbb{R}). \tag{7.40}$$

For every integer $K \ge 0$,

$$\int_{|t|>L} |\widehat{\Phi}(t)|\, dt \le L^{-K} \int_{\mathbb{R}} |t|^K |\widehat{\Phi}(t)|\, dt.$$

The last integral is finite. Therefore the tail in (7.39) is smaller than any prescribed negative power of L after choosing K. This choice follows the fixed kernel and J. Only its constant and the final threshold depend on K.

The low-frequency bound (7.35) is L^{-26}. Its integral against $|\widehat{\Phi}|$, and the tail just estimated, give

$$\left| \int_{\mathbb{R}} \widehat{\Phi}(t)\mathcal{U}_{0,X}(-2\pi t)\, dt \right| \le L^{-24} \tag{7.41}$$

for every active soft source scale.

Choose ψ to be a smooth dyadic partition. More explicitly, let η be smooth, equal to one on $(-\infty, 1]$, equal to zero on $[2, \infty)$, and valued in $[0,1]$, and set $\psi(v) = \eta(v) - \eta(2v)$. For $2 \le x \le 2^K$,

$$\sum_{j=1}^{K} \psi(x/2^j) = \eta(x/2^K) - \eta(x) = 1. \tag{7.42}$$

Take $K = 1 + \lceil 2L/\log 2 \rceil$ and $X_j = 2^j$. Every shared product satisfies $D \le \exp(L^{.60}) \le \exp(L/20)$ eventually. On the support of $\chi(n/N)$ this puts Dn in $[2, e^{2L}]$ eventually, so (7.42) applies. If the jth summand is nonzero,

$$\frac{N}{8} < X_j < \frac{3}{2} N \exp(L/20).$$

Thus every contributing scale lies in $[\exp(.9L), \exp(1.1L)]$ for large L. At each such scale the two hard source scales lie in the moment range. Every other scale gives zero in (7.39) after inversion.

Sum (7.39) over j. The dyadic partition is applied for each shared tuple, after which $\sum_{\mathbf{s}} \mu_{\mathrm{sh}}(\mathbf{s}) = 1$. Since $K \le L^2$ eventually, (7.41) proves the following statement.

Proposition 7.3 (Centered marked correlation)**.** *Fix $0 < \varepsilon < 1/4$ and a real smooth cutoff χ supported in $(1/4, 3/4)$. Choose the graph construction at the precision above, and then fix J at least its prescribed lower bound. For all sufficiently large N which are N^{ε}-rough, every real function H on the positive integers satisfying*

$$|H(n)| \le 1, \qquad H(dn) = H(n) \quad (d \text{ positive and } Q\text{-smooth},\ n \ge 1)$$

satisfies

$$\left| \frac{1}{N} \sum_{0<n<N} \chi(n/N) W(n) W(N-n) F_0(n) H(N-n) \right| \le L^{-22}. \tag{7.43}$$

The threshold is uniform over these functions H. The construction parameters precede J, and J precedes the final threshold in N.

Proof. Sum (7.39) over the $K \le L^2$ dyads in (7.42) and apply (7.41). All endpoint bounds used above depend only on $|H| \le 1$ and its stated invariance, so the threshold is uniform in H. □

The two applications use the smoothness indicator and the invariant principal model for H. Their common weight is $\chi(n/N)W(n)W(N-n)/N$. The positive baseline and the final implication from these centered estimates are given in Proposition 8.5.

8. Positive correlation and smooth summands

Fix $0 < \varepsilon < 1/4$, put $L = \log N$ and $y = N^{\varepsilon}$, and use the prime groups $\mathcal{P}_g$ of Section 5. Throughout this section

$$q = \frac{1}{256}, \qquad d_0 = \frac{49}{100}, \qquad Q = \exp(L^{d_0}).$$

Write $s = |\mathcal{S}| + |\mathcal{B}|$, and retain the notation $V_g = \sum_{p \in \mathcal{P}_g} p^{-1}$ and $\mu_g(p) = (V_g p)^{-1}$. For positive integers n, define the original one-mark weights by

$$\omega_g(n) = |\{p \in \mathcal{P}_g : p \mid n\}|, \qquad w_g(n) = \frac{\omega_g(n) q^{\omega_g(n)-1}}{V_g}, \qquad W(n) = \prod_g w_g(n). \tag{8.1}$$

The value of $w_g(n)$ is zero when $\omega_g(n) = 0$. In particular all these weights are nonnegative. The group estimates (5.1) imply, for large L,

$$s \le 2 \log L, \qquad \frac{1}{2} \le V_g \le \frac{3}{2}, \qquad e^{L^{27/100}} < p \le e^{L^{46/100}} \quad (p \in \mathcal{P}_g). \tag{8.2}$$

These bounds involve no slot parameter J.

8.1. The original marked mass.

Lemma 8.1 (Uniform marked CRT interval bound)**.** *Suppose that N is coprime to every group prime. For sufficiently large L, uniformly over integer intervals $I \subseteq [1, N)$ with $|I| \ge e^{L/2}/2$,*

$$\sum_{n \in I} W(n) W(N-n) \ge |I| L^{-15}. \tag{8.3}$$

Proof. We give the interval comparison with its original harmonic normalizations. Expand each factor in (8.1) by choosing a marked prime $p_g \mid n$ and a marked prime $r_g \mid N-n$. The choice $p_g = r_g$ contributes zero, since $p_g \nmid N$. For a fixed choice with $p_g \ne r_g$ in every group, put

$$A = \prod_g p_g r_g, \qquad \mathcal{P}' = \left(\bigcup_g \mathcal{P}_g \right) \setminus \{p_g, r_g : g\}.$$

The marked congruences specify a single residue $z \pmod A$. Their contribution before the factors V_g^{-2} are inserted is

$$\sum_{n\in I} \mathbf{1}_{n\equiv z \bmod A} \prod_{p\in\mathcal{P}'} \Big(1-(1-q)(\mathbf{1}_{p|n}+\mathbf{1}_{p|N-n})\Big). \tag{8.4}$$

The two residue classes in each parenthesis are distinct. Thus its loss lies in $[0,1]$, and its uniform residue mean is

$$\delta_p = \frac{2(1-q)}{p}.$$

Set $\mu = \prod_{p\in\mathcal{P}'}(1-\delta_p)$. We claim that (8.4) is at least

$$\frac{|I|}{2A}\mu. \tag{8.5}$$

Let K be the least odd integer at least $40\log L$. The odd Bonferroni truncation of the product in (8.4) is a pointwise lower bound. For a subset of j unmarked primes, the Chinese remainder theorem gives 2^j distinct residue classes modulo A times their product. Each class contains $|I|/(A\prod p)$ integers of I, with error at most one. Consequently the sum of the truncated product differs from

$$\frac{|I|}{A}\sum_{j=0}^{K}(-1)^j e_j((\delta_p)_{p\in\mathcal{P}'})$$

by at most the following quantity, where e_j denotes the elementary symmetric polynomial of degree j,

$$B_K = \sum_{j=0}^{K} 2^j \binom{|\mathcal{P}'|}{j} \le (K+1)\big(2(|\mathcal{P}'|+1)\big)^K. \tag{8.6}$$

The coefficients $(1-q)^j$ only decrease this error.

By (8.2),

$$\sum_{p\in\mathcal{P}'}\delta_p \le 2\sum_g V_g \le 3s \le 6\log L.$$

Writing $e=\exp(1)$, the discarded mean terms satisfy

$$\begin{aligned}\sum_{j>K} e_j((\delta_p)_p) &\le e^{-K}\prod_{p\in\mathcal{P}'}(1+e\delta_p)\\ &\le \exp\left(-K+e\sum_{p\in\mathcal{P}'}\delta_p\right) \le L^{-22}.\end{aligned}$$

For large L, every $\delta_p \le 1/2$. The inequality $\log(1-t)\ge -2t$ for $0\le t\le 1/2$ gives

$$\mu \ge \exp\left(-2\sum_{p\in\mathcal{P}'}\delta_p\right) \ge L^{-12}. \tag{8.7}$$

Also

$$\log A \le 2sL^{46/100} \le 4L^{46/100}\log L, \qquad |\mathcal{P}'| \le e^{L^{46/100}},$$

so (8.6) gives $B_K \le \exp(O(L^{46/100}\log L))$. The lower truncation therefore bounds (8.4) from below by

$$\frac{|I|}{A}(\mu - L^{-22}) - B_K.$$

Since $|I| \geq e^{L/2}/2$, it follows uniformly in the marked primes, the target residue and the interval that

$$L^{-22} \leq \frac{\mu}{4}, \qquad B_K \leq \frac{|I|\mu}{4A}$$

for all sufficiently large L. The lower Bonferroni bound, the mean tail estimate and these two inequalities prove (8.5).

Sum (8.5) over the marked choices with their original coefficient $\prod_g V_g^{-2}$. The mean factors across groups. Its factor for a group g is

$$\beta_g = \sum_{\substack{p,r \in \mathcal{P}_g \\ p \neq r}} \mu_g(p)\mu_g(r) \prod_{\ell \in \mathcal{P}_g \setminus \{p,r\}} \left(1 - \frac{2(1-q)}{\ell}\right). \tag{8.8}$$

The same logarithmic inequality shows that the product in this sum is at least e^{-4V_g}. Moreover,

$$\sum_{\substack{p,r \in \mathcal{P}_g \\ p \neq r}} \mu_g(p)\mu_g(r) = 1 - \sum_{p \in \mathcal{P}_g} \mu_g(p)^2 \geq 1 - \frac{1}{V_g \min \mathcal{P}_g} \geq \frac{1}{2}$$

eventually. Hence $\beta_g \geq \frac{1}{2}e^{-4V_g}$, and therefore

$$\sum_{n \in I} W(n)W(N-n) \geq \frac{|I|}{2} \prod_g \beta_g \geq |I|\, 2^{-(s+1)} \exp\left(-4 \sum_g V_g\right). \tag{8.9}$$

Here the reciprocal factor A^{-1} in the CRT mean has supplied one factor p^{-1} for each marked prime. In combination with V_g^{-2}, these are exactly the harmonic probabilities in (8.8).

Finally, $\sum_g V_g \leq 3s/2$, $s \leq 2 \log L$ and $L \geq e$ give

$$2^{-(s+1)} \exp\left(-4 \sum_g V_g\right) \geq e^{-7s-1} \geq e^{-14 \log L - 1} \geq L^{-15}.$$

This proves (8.3). □

Let χ be the fixed smooth cutoff used in the comparison section. We use

$$\begin{gathered} 0 \leq \chi \leq 1, \qquad \chi(1-t) = \chi(t), \\ \chi(t) = 0 \quad (t \notin (1/4, 3/4)), \qquad \chi(t) = 1 \quad (1/3 \leq t \leq 2/3). \end{gathered} \tag{8.10}$$

Lemma 8.2. *For every sufficiently large N that is N^{ε}-rough,*

$$\sum_{0<n<N} \chi(n/N)W(n)W(N-n) \geq \frac{N}{4} L^{-15}. \tag{8.11}$$

Proof. Every group prime is at most $\exp(L^{46/100}) < N^{\varepsilon}$ for large N, so roughness gives $p \nmid N$ for every such prime. Put $h = \lfloor N/3 \rfloor$ and

$$I = [h+1, 2h) \cap \mathbb{Z}.$$

For $N \geq 24$, this interval lies in $[N/3, 2N/3]$, has $|I| = h - 1 \geq N/4$, and consequently has $|I| \geq \sqrt{N}/2 = e^{L/2}/2$. Apply Lemma 8.1. On I the cutoff equals one, while all summands outside I are nonnegative. This proves (8.11). □

8.2. **The principal baseline.** Use the extension $\mathcal{M}$ from (4.14) and put

$$m(n) = \mathcal{M}(n^\sharp), \qquad F_0(n) = f_y(n) - m(n) \quad (n \ge 1), \qquad c_\varepsilon = \frac{1}{2}\rho(2/\varepsilon) > 0. \tag{8.12}$$

Since $Q \le y$ for large N, this F_0 is the centered function F in (4.14).

Lemma 8.3. *For sufficiently large N,*

$$c_\varepsilon \le \mathcal{M}(x) \le 1 \qquad (x > 0), \qquad c_\varepsilon \le m(n) \le 1 \qquad (n \ge 1).$$

Both f_y and m are invariant under multiplication by a positive Q-smooth integer.

Proof. For $e^{L/2} \le x \le e^{2L}$, Lemma 4.2 gives

$$\mathcal{M}(x) = \rho\Big(\frac{\log x}{\varepsilon L}\Big) + O_\varepsilon(L^{-1/50}(\log L)^2).$$

The argument of ρ is at most $2/\varepsilon$, and ρ is nonincreasing. The error is eventually at most $\rho(2/\varepsilon)/2$, so $\mathcal{M}(x) \ge c_\varepsilon$. The upper bound on this interval is also part of Lemma 4.2. Outside the interval, $\mathcal{M}(x) = \rho(1/\varepsilon)$, which belongs to $[\rho(2/\varepsilon), 1]$. This proves both displayed bounds.

If D is Q-smooth, then $(Dn)^\sharp = n^\sharp$, so $m(Dn) = m(n)$. Since every prime factor of D is at most $Q \le y$, the integer Dn is y-smooth exactly when n is y-smooth. Thus $f_y(Dn) = f_y(n)$ as well. □

For real functions F, G on the positive integers, write

$$\mathcal{C}_N(F,G) = \frac{1}{N}\sum_{0<n<N} \chi(n/N)W(n)W(N-n)F(n)G(N-n). \tag{8.13}$$

Lemma 8.4. *For every sufficiently large N that is N^ε-rough,*

$$\mathcal{C}_N(m,m) \ge \frac{c_\varepsilon^2}{4}L^{-15}. \tag{8.14}$$

Proof. Each factor $m(n)m(N-n)$ in (8.13) is at least c_ε^2. The remaining weight is nonnegative. Apply Lemma 8.2 and divide by N. □

8.3. **The fixed-exponent implication.** The universal raw-correlation estimate supplied by the comparison section, (7.43), has the following precise form. For every fixed $0 < \varepsilon < 1/4$ and all sufficiently large N^ε-rough targets N,

$$|\mathcal{C}_N(F_0, H)| \le L^{-22} \tag{8.15}$$

for every real function H on the positive integers satisfying

$$|H(n)| \le 1, \qquad H(Dn) = H(n) \quad (n \ge 1,\ D \text{ a positive } Q\text{-smooth integer}). \tag{8.16}$$

The function H may be chosen after N and its prime groups. The graph construction and its constants are fixed before that choice.

Proposition 8.5 (Fixed-exponent implication)**.** *Fix $0 < \varepsilon < 1/4$ and assume (8.15) for the functions (8.16). Every sufficiently large N^ε-rough integer N is a sum of two positive N^ε-smooth integers.*

Proof. The functions $H = f_y$ and $H = m$ satisfy (8.16) by Lemma 8.3. Thus

$$|\mathcal{C}_N(F_0, f_y)| \le L^{-22}, \qquad |\mathcal{C}_N(F_0, m)| \le L^{-22}. \tag{8.17}$$

The identity

$$f_y(n)f_y(N-n) - m(n)m(N-n) = F_0(n)f_y(N-n) + m(n)F_0(N-n)$$

is exact. Under $n \mapsto N-n$, the interval $0 < n < N$ and the weight $\chi(n/N)W(n)W(N-n)$ are unchanged. Consequently

$$\mathcal{C}_N(f_y, f_y) = \mathcal{C}_N(m, m) + \mathcal{C}_N(F_0, f_y) + \mathcal{C}_N(F_0, m). \tag{8.18}$$

This reflection uses the same normalization at both endpoints. Equivalently, for the reciprocal cutoff $\Psi(t) = t\chi(t)$,

$$\frac{\Psi(n/N)}{n} = \frac{\chi(n/N)}{N}, \qquad \frac{t}{1-t}\Psi(1-t) = \Psi(t) \quad (0 < t < 1).$$

Since $c_\varepsilon > 0$ is fixed and $L^{-7} \longrightarrow 0$,

$$2L^{-22} < \frac{c_\varepsilon^2}{4} L^{-15}$$

for all sufficiently large N. Combining (8.14), (8.17) and (8.18) yields

$$\mathcal{C}_N(f_y, f_y) > 0. \tag{8.19}$$

Every summand in this last correlation is nonnegative. Hence some integer n with $0 < n < N$ satisfies $f_y(n)f_y(N-n) = 1$. The two positive integers n and $N-n$ are therefore y-smooth and sum to N, as required. □

Completion of the proof.

Proof of Theorem 1.1. Fix $0 < \varepsilon < 1/4$ and choose the smooth symmetric cutoff used in Section 8. Proposition 7.3 supplies (8.15) uniformly for all the endpoint functions in (8.16). Proposition 8.5 therefore gives the required representation for every sufficiently large N^ε-rough target N. Lemma 2.2 then gives the assertion for every fixed positive exponent and every sufficiently large target. □

Acknowledgments

AI systems developed by OpenAI assisted with mathematical exploration, proof development, and the preparation of this manuscript.

References

[1] A. Balog, *On additive representation of integers*, Acta Math. Hungar. **54** (1989), 297–301. doi:10.1007/BF01952060.
[2] T. F. Bloom, *Erdős Problem #334*, Erdős Problems, electronic problem collection. Problem page.
[3] H. Davenport, *Multiplicative Number Theory*, third edition, revised by H. L. Montgomery, Graduate Texts in Mathematics, vol. 74, Springer, New York, 2000.
[4] P. Erdős, *Problems and results on consecutive integers*, Publ. Math. Debrecen **23** (1976), 271–282. Publisher's record.
[5] P. Erdős, *Some new problems and results in number theory*, in *Number Theory* (Mysore, 1981), Lecture Notes in Mathematics, vol. 938, Springer, Berlin, 1982, 50–74. Electronic manuscript.
[6] B. Green, *100 open problems*, manuscript, Problem 59. Electronic manuscript.
[7] OpenAI, *The Quasi-Riemann Hypothesis. A Zero-Free Half-Plane $\Re s > 7/8$*, preprint, September 30, 2026. Electronic manuscript.
[8] OpenAI, *Weighted dilation graphs, smooth shifted primes and totient fibers*, preprint, September 24, 2026. Electronic manuscript.