hexagon:2611.01210v1

Reflected Liouville correlations and joint smoothness laws

Brice Pouly

With substantial assistance from OpenAI language models.

Abstract

Does each prescribed pair of Liouville signs occur in a quarter of the ordered decompositions of every large integer? We reduce this conjecture of Corrádi and Kátai to a marked correlation estimate and give a candidate proof of that estimate. Random colours on prime divisors make the two marking weights independent on coprime summands. Their exact mean can be inverted on a set of density tending to one. With a suitable number of colours, a marked saving $(\log N)^{-\beta}$ for any fixed $\beta > 51/160$ suffices for the reduction.

At one summand we prove, using the classical prime number theorem, that the sign is fair even after imposing smoothness and testing the part supported on small primes. At two summands the marked estimate gives the product of two fair signs and two Dickman smoothness laws, with uniform absolute error $o(N)$ over intervals and over both threshold exponents in $[0, 1/4]$. The candidate proof adapts the dilation graphs of Howe, Huffman, Kohlmeyer and Kruer to signed sources and dispersed prime groups. It uses a uniform estimate for primes twisted by Dirichlet characters, deduced in that work from a zero-free half-plane $\Re s > 7/8$.

Contents

1 The reflected problem and the main conclusions 2

2 Prime observations and the marked estimate 6

3 Averaging the marking weights 8

4 Recovering ordinary reflected averages 13

5 One-variable cancellation and the two centres 16

6 A parametric transfer from the ideal graph 20

7 From the physical moment to the marked correlation 35

8 Recovering the joint laws and their local variants 48

9 Range of the conclusions 55

References 57

1 The reflected problem and the main conclusions

Every representation $N = n + (N - n)$ carries two signs of the Liouville function

$$\lambda(n) = (-1)^{\Omega(n)},$$

where Ω counts prime factors with multiplicity. The prime number theorem gives equal proportions of positive and negative signs at either summand separately. It does not determine whether the two signs agree. The reflected correlation

$$C_\lambda(N) = \sum_{1 \le n < N} \lambda(n)\lambda(N-n) \tag{1.1}$$

measures exactly this missing information.

For $\sigma, \tau \in \{-1, 1\}$ write

$$R_{\sigma,\tau}(N) = \#\{1 \le n < N : \lambda(n) = \sigma,\ \lambda(N-n) = \tau\}.$$

These are ordered representations. Expanding the two sign indicators gives

$$4R_{\sigma,\tau}(N) = N - 1 + (\sigma + \tau) \sum_{n<N} \lambda(n) + \sigma\tau C_\lambda(N). \tag{1.2}$$

Since the middle sum is $o(N)$, the assertion $C_\lambda(N) = o(N)$ is equivalent to four frequencies $1/4 + o(1)$. We use "the reflected Corrádi–Kátai conjecture" for this assertion, with N tending to infinity through all positive integers; see [1, 2, 5]. There is no average over the total and no logarithmic weight on the summation variable.

1.1 Main statements

The argument isolates a marked correlation estimate, stated in Proposition 2.2, in which each summand is weighted by its prime divisors in two short-power ranges. The source may retain a restricted part of its factorisation, and the target may retain the complementary factors and prime counts. The statement extends the centred smoothness correlation in the preprint [7]. Its precise information restrictions matter: revealing the whole source would make the assertion false.

The first two theorems are implications. Their common arithmetic premise has a candidate proof in Sections 6 and 7. Proposition 5.4 gives the corresponding one-summand law from classical estimates alone.

Theorem 1.1 (Reflected signs from a marked estimate)**.** *Assume the F_λ case of (2.11) for one fixed $0 < \varepsilon < 1/4$ with the uniformity stated there. Then*

$$C_\lambda(N) = o(N) \qquad (N \longrightarrow \infty) \tag{1.3}$$

along all integers. For intervals $I \subseteq (0,1)$, with the full interval $(0,1)$ allowed,

$$\#\{1 \le n < N : n/N \in I,\ \lambda(n) = \sigma, \lambda(N-n) = \tau\} = \frac{|I|}{4} N + o(N). \tag{1.4}$$

The error is uniform over all intervals I and the four sign pairs. The assertion holds for both parities of N.

The extra source terms in the marked statement retain smoothness at the same time as the sign. To formulate the resulting law, fix $0 < \varepsilon < 1/4$ and put

$$S_N(v) = \mathbf{1}_{P^+(v) \le N^\varepsilon}, \qquad r = \rho(1/\varepsilon), \qquad \pi_1 = r, \quad \pi_0 = 1 - r.$$

Here $P^+(1) = 1$ and ρ is the Dickman function, defined by $\rho(u) = 1$ for $0 \le u \le 1$, $\rho(u) = 0$ for $u < 0$, and $u\rho'(u) = -\rho(u-1)$ for $u > 1$. For its relation to smooth-number densities, see [11].

Theorem 1.2 (The joint sign and smoothness law)**.** *Assume the three cases of (2.11) for the fixed ε above. For intervals $I \subseteq (0,1)$, $i, j \in \{0,1\}$ and $\sigma, \tau \in \{-1,1\}$,*

$$\#\{1 \le n < N : n/N \in I,\ S_N(n) = i,\ S_N(N-n) = j,\ \lambda(n) = \sigma, \lambda(N-n) = \tau\} \\ = \frac{|I|\pi_i\pi_j}{4} N + o_\varepsilon(N). \quad (1.5)$$

The error is uniform over all intervals I and all sixteen states. Thus the limiting four-coordinate law is the product of two Bernoulli variables of parameter r and two fair signs. In particular, each sign pair occurs among the two N^ε-smooth summands with count $|I|r^2N/4 + o_\varepsilon(N)$, uniformly in I.

Corollary 8.7 allows different exponents at the two summands and gives a uniform absolute error through a finite partition of the threshold range. In the analytic estimates the exponent ε is fixed before the limit. The interval may depend on N, because the error is absolute and normalized by N. No relative error $o(N|I|)$ is asserted for shrinking intervals, and the proof gives no relative estimate for the rare population as ε tends to zero. Section 8 gives two further consequences of the same local estimates. Nonzero Möbius signs introduce the factor

$$\mathfrak{S}_2(N) = \prod_{p^2 \mid N} (1 - p^{-2}) \prod_{p^2 \nmid N} (1 - 2p^{-2}), \quad (1.6)$$

and the sign conclusions extend uniformly to two possibly different multiplicative functions $f, g : \mathbb{N} \to \{-1, 1\}$ satisfying

$$\sum_{p : f(p) \ne -1} p^{-\beta} \le K, \qquad \sum_{p : g(p) \ne -1} p^{-\beta} \le K, \qquad \beta < 1, \quad (1.7)$$

with β, K fixed. Values at prime powers are unrestricted, and f, g may vary with N. These extensions require the local-test content of the transfer, not just (1.3).

1.2 Why random prime colours remove a weight

The dilation-graph argument of [7] uses a weight that vanishes unless an integer has a prime factor in every prescribed group. A fixed collection of groups can leave a sparse subset of integers. We instead distribute each prime independently among several colours. For an integer with k distinct prime factors in a block and s colours, let K_j count its factors of colour j. If $0 < \theta < 1$ is fixed, then

$$\mathbb{E} \prod_{j=1}^{s} K_j \theta^{K_j - 1} = \frac{(k)_s}{s^s} \theta^{k-s}, \qquad (k)_s = k(k-1)\cdots(k-s+1), \quad (1.8)$$

with value zero for $k < s$. Each factor K_j selects one distinct marked prime of colour j; the remaining $k - s$ factors supply the power of θ.

For example, with three distinct prime factors and two colours, six of the eight assignments occupy both colours. On each of these assignments the product is 2θ, so its expectation is $3\theta/2$, exactly as in (1.8). If two integers are coprime, their sets of prime factors are disjoint and these colouring expectations multiply. This is the point at which the arithmetic condition on the two summands is used.

Proposition 3.5 shows that, for a suitable number of colours, a marked saving $(\log N)^{-\beta}$ with any fixed $\beta > 51/160$ is enough for ordinary cancellation. The inverse of (1.8) is polynomial in $\log N$ when the number of factors is $O(\log\log N)$ and at least the number of colours. Both requirements fail on a set of density tending to zero. Consequently a sufficiently accurate marked correlation controls the unweighted correlation after averaging over colours. The randomness belongs only to this proof device; the function λ and the integers being counted remain fixed.

The subsequent common-divisor decomposition is equally concrete. If $d = (n, N)$, $n = du$ and $N = dT$, then $(u, T) = 1$ and

$$\lambda(n)\lambda(N-n) = \lambda(u)\lambda(T-u).$$

It allows the coprime calculation to address every representation. A uniform version with a common outer scale also preserves the smoothness cutoff and the small-prime tests.

1.3 Relation to earlier results

Mangerel's nonextremality theorem gives $|C_\lambda(N)| < N - 1$ for $N \geq 11$ [2, Theorem 1.2]. This ensures both products of signs; it does not by itself distinguish the two equal-sign patterns. Meng establishes all four patterns for every $N \geq 2$ outside the sharp exceptional set $\{2, 3, 4, 5, 6, 9, 10\}$ [4, Theorem A]. These existence assertions retain information at small totals that an asymptotic theorem does not provide.

The author's *Reflected sign patterns of the Liouville function at every large total* gives $R_{\sigma,\tau}(N) \geq c_\eta N^{1-\eta}$ for each fixed $\eta > 0$ and sufficiently large N, together with a uniform prime-budget extension and separate logarithmic bounds at prime totals [6, Theorems 1.1, 1.2 and 1.4]. The fixed-budget uniformity and the possibility that the function varies with the total are already present there. Its reflection-to-dilation argument studies near-extremality and reconstructs a real character. The marked estimate gives asymptotic frequencies, allows two possibly different multiplicative functions with unrestricted prime-power values, and retains the joint smoothness information. The earlier character-rigidity results are not used to prove the transfer.

De Koninck, Germán and Kátai obtain conditional bounds for the reflected Liouville convolution from sufficiently close Siegel zeros, on ranges tied to the associated conductor and with a coprimality restriction [12]. Under GRH, Mangerel obtains quantitative lower bounds on specified domains, including $p\exp(-C(\log\log p)^6)$ for the four patterns at large prime totals [3, Proposition 1.4]. These hypotheses and domains are part of their statements.

Two other types of averaging are relevant. Cantarini, Gambini and Zaccagnini derive explicit formulae for weighted averages of the Liouville and Möbius convolutions, under RH and simplicity of the nontrivial zeta zeros [13, Theorem 1 and Section 4]. Krishnamoorthy proves that, for every $A > 0$,

$$\#\{N \leq X : |C_\lambda(N)| > \delta N\} \ll_A \delta^{-2} X(\log X)^{-A}$$

in the range of his theorem [5, Theorem 2]. The conclusion sought in (1.3) concerns every sufficiently large total. Slothouber's *Phase cones for Gamma-weighted atoms of additive Dirichlet series* studies sums of residues at a candidate singularity. Its Liouville application

requires additional control of the phases or of gaps between zero pairs [24, Section 4.1]. Section 9 gives an explicit multiplicative example in which a small Cesàro average coexists with maximal reflected correlations on a subsequence.

Ordinary correlation estimates at fixed affine forms require a separate comparison. The Liouville theorem in [20] fixes $a_1, a_2 \geq 1$ and $b_1, b_2 \geq 0$, with $a_1 b_2 - a_2 b_1 \neq 0$, and bounds the ordinary correlation at $a_1 n + b_1$ and $a_2 n + b_2$. Its constants may depend on those forms, and it states no uniformity for coefficients growing with the summation scale. The pair $n, N - n$ has a negative second slope and an intercept varying with the total. The fixed-form theorem therefore does not imply the reflected assertion considered here.

The author's *Joint cofactor signs in prime-factor windows on consecutive integers* counts signs of cofactors of v and $v-1$ with specified large prime divisors [25]. Here the two summands add to a prescribed total, and the target can observe all factors outside the coloured groups.

Joint smoothness laws at a fixed shift also have a distinct history. Teräväinen proves the product Dickman law for $n, n+1$ in logarithmic density [18, Theorem 1.14]; the preprint [19] gives the corresponding ordinary-density law. Both statements keep the shift fixed. Theorem 1.2 concerns instead two summands of each prescribed total and includes their signs.

For additive representations by smooth integers, Balog's work [21] is an earlier antecedent; its exponent $4/(9\sqrt{e})$ is recalled in the introduction to [22]. Jain proves that, for each $\beta > 103/(256\sqrt{e})$, every sufficiently large N is a sum $a + b$ with $P^+(a) \leq N^\beta$ and $P^+(b) \leq (\log N)^{K(\beta)}$ [22, Theorem 1]. The polylogarithmic cutoff is a different regime from the fixed-power cutoff in Theorem 1.2; a positive representation count at that cutoff would require finer control than the absolute error in Corollary 8.7.

Howe, Huffman, Kohlmeyer and Kruer prove in their preprint that every sufficiently large integer is a sum of two N^ε-smooth positive integers, for each fixed $\varepsilon > 0$ [7]. Their construction provides the dilation graphs, character estimates and invariant smoothness centre used here. It imports the zero-free half-plane $\Re s > 7/8$ for Dirichlet L-functions from [8], and develops the weighted-dilation methods of [9]. The added tasks are to admit dispersed prime groups, preserve signed observations under the graph dilations, and remove the marking weights. We identify these tasks at their actual points of use.

For the local variants, the unsigned asymptotic for two squarefree summands and its local factor are classical; see [15, 16] and the account in [17, pp. 591–592]. The additional assertion in Proposition 8.3 is the distribution of the signs within that population, including the joint smoothness restriction. Tian and Yang obtain lower bounds at every sufficiently large total and quarter frequencies at almost all totals for the unitary Möbius function $(-1)^{\omega(n)}$ [14, Theorems 1.5 and 1.7]. This function satisfies (1.7) with budget zero. Their exact finite exceptions, given in Theorem 1.6, remain separate information. Their Conjecture 4.3 concerns the ordinary Möbius reflected correlation; it follows from Proposition 8.3 by summing the four nonzero sign counts with coefficient $\sigma\tau$.

1.4 How the estimates fit together

Fix the total and remove its common divisors with the first summand. On the remaining primitive pairs the colours of primes at the two sites are independent. Sections 3 and 4 turn this identity into the implication from the marked estimate to ordinary counting. The one-variable centres in Section 5 use the classical prime number theorem and also give the one-site sign law with local tests.

The step specific to the two-site conjecture is Proposition 7.1. Section 6 passes from an ideal operator to a moment on integer positions. Section 7 applies that moment to two separately specified functions and recovers counting measure. The uniform twisted-prime estimate (6.22)

enters the comparison at nonprincipal characters. Section 8 derives the joint laws from the resulting marked input.

2 Prime observations and the marked estimate

We first specify what may be observed at the two sites. This prevents a bounded test from carrying the very sign whose cancellation is being sought.

2.1 One outer scale and two prime ranges

Let N be the outer total, put $L = \log N$ and $\ell = \log L$, and fix $0 < \varepsilon < 1/4$. Throughout the proof,

$$B = \exp(L^{.26}), \quad Q = \exp(L^{.49}), \quad y = \exp(\varepsilon L), \quad G = \exp(L^{.20}). \tag{2.1}$$

The letters B, Q, y always refer to this outer scale. Auxiliary integer totals T are allowed in the range

$$e^{L-L^{.20}} \le T \le e^{L}. \tag{2.2}$$

The two prime blocks are

$$\mathcal{B}_i(T) = \{p : e^{L^{a_i}} < p \le e^{L^{b_i}},\ p \nmid T\}, \quad (a_1, b_1) = (.27, .36), \quad (a_2, b_2) = (.39, .46). \tag{2.3}$$

Write $w_i = b_i - a_i$ and

$$s_i = \lfloor 3w_i\ell/4 \rfloor, \qquad s = s_1 + s_2 \le .12\ell.$$

All assertions concern N large enough that $s_1, s_2 \ge 1$ and $G < B < \min \mathcal{B}_1(T) < Q < y$. Each prime in block i is independently assigned one of s_i colours. A realization $\mathfrak{c}$ gives pairwise disjoint groups $\mathcal{P}_g$, with harmonic masses

$$V_g = \sum_{p \in \mathcal{P}_g} p^{-1}.$$

The block totals and the set $\mathcal{P} = \mathcal{B}_1(T) \cup \mathcal{B}_2(T)$ do not depend on the colouring.

For an integer $v \ge 1$, define

$$v^{\circ} = \prod_{p \notin \mathcal{P}} p^{v_p(v)}, \qquad v^{\flat} = \prod_{p \le Q} p^{v_p(v)}, \qquad v^{\sharp} = \prod_{p > Q} p^{v_p(v)},$$
$$\omega_g(v) = \#\{p \in \mathcal{P}_g : p \mid v\}, \qquad \Omega_g(v) = \sum_{p \in \mathcal{P}_g} v_p(v).$$

The multiplicity counts Ω_g are needed under dilation; the weights themselves use the distinct-prime counts ω_g.

2.2 Different information at the source and target

For each colouring use the two observations

$$Y_{\mathrm{src}}(v) = \big((v^{\circ})_{\le Q}, (\Omega_g(v))_g\big), \qquad Y_{\mathrm{tgt}}(v) = \big(v^{\circ}, (\Omega_g(v))_g\big). \tag{2.4}$$

Here $u_{\le z} = \prod_{p \le z} p^{v_p(u)}$. The target observation determines both

$$\lambda(v) = \lambda(v^{\circ})(-1)^{\sum_g \Omega_g(v)} \quad \text{and} \quad S(v) = \mathbf{1}_{P^+(v) \le y}.$$

The second assertion holds because every prime in a block is below y. The source observation does not expose the factors above Q. It also withholds the identities of the block primes; it retains only their group counts.

In particular, an arbitrary function of $v_{\le Q}$ is not an admissible source test: it could read the identity of a selected block prime. An arbitrary function of $v_{\le B}$ is admissible, since B is below both blocks. These are precisely the smaller local factors needed in Section 8.

Example 2.1 (What remains after a parity correction)**.** The source test $(-1)^{\sum_g \Omega_g(v)}$ cancels the parity of the block factors in $\lambda(v)$. Multiplying also by $\lambda((v^\circ)_{\le Q})$ leaves $\lambda(v^\sharp)$. This last sign is still unobserved at the source. Permitting the full exterior v° at the source would instead allow the multiplier $\lambda(v)$ and destroy cancellation. The asymmetry in (2.4) is therefore a substantive restriction.

Set $\theta = 1/256$ and define

$$W_{\mathfrak{c}}(v) = \prod_g \frac{\omega_g(v)\theta^{\omega_g(v)-1}}{V_g}, \tag{2.5}$$

with a factor equal to zero when $\omega_g(v) = 0$. We use this normalized weight only when all $V_g > 0$. Its unnormalized counterpart

$$U_{\mathfrak{c}}(v) = \prod_g \omega_g(v)\theta^{\omega_g(v)-1}$$

is defined for every colouring and satisfies $0 \le U_{\mathfrak{c}}(v) \le 1$.

2.3 The three centred sources

Section 5 defines bounded functions $\mathcal{M}_0, \mathcal{M}_-$ with the following properties. For $r = \rho(1/\varepsilon)$,

$$\mathcal{M}_0(x) = \rho\Big(\frac{\log x}{\log y}\Big) + O_\varepsilon(L^{-.02}\ell^2) \qquad (e^{L/2} \le x \le e^{2L}), \tag{2.6}$$

$$\|\mathcal{M}_-\|_\infty \ll_\varepsilon L^{-.49}. \tag{2.7}$$

Outside that interval $\mathcal{M}_0$ is extended by r and $\mathcal{M}_-$ by zero. The sources in the marked estimate are

$$F_\lambda(v) = \lambda(v), \qquad F_0(v) = S(v) - \mathcal{M}_0(v^\sharp), \qquad F_*(v) = \lambda(v)S(v) - \lambda(v^\flat)\mathcal{M}_-(v^\sharp). \tag{2.8}$$

They are bounded uniformly for large N. If d is Q-smooth, then

$$F_0(dv) = F_0(v), \qquad F_\lambda(dv) = \lambda(d)F_\lambda(v), \qquad F_*(dv) = \lambda(d)F_*(v). \tag{2.9}$$

These identities hold with multiplicity and are the reason for retaining the centres until after the weights have been removed.

2.4 A uniform statement for the modified groups

We call a colouring regular if $1/2 \le V_g \le 3/2$ for every group and if, for every dyadic interval $[u, 2u]$ meeting its block, every subinterval $J \subseteq [u, 2u]$, every Dirichlet character ξ of modulus at most e^{2L}, and every real $|t| \le 2e^{3L}$,

$$\Big|\sum_{p \in J \cap \mathcal{B}_i(T)} (\mathbf{1}_{p \in \mathcal{P}_g} - s_i^{-1})\xi(p)p^{it}\Big| \le u\exp(-L^{.27}/100). \tag{2.10}$$

Increasing the absolute constant on the right, if necessary, gives the same conclusions. Lemma 3.1 proves that regular colourings have probability tending to one much faster than any inverse power of L.

Proposition 2.2 (Marked transfer at a common outer scale, candidate)**.** *For each fixed $0 < \varepsilon < 1/4$ and real $\chi \in C_c^\infty((0,1))$, there are $C_{\varepsilon,\chi}$ and L_0 with the following property. For $L \geq L_0$, every integer T satisfying (2.2), every regular colouring of (2.3), and every pair of real bounded functions a, b on the observations in (2.4),*

$$\frac{1}{T}\left|\sum_{\substack{1\leq u<T\\(u,T)=1}} \chi(u/T)W_{\mathfrak{c}}(u)W_{\mathfrak{c}}(T-u)F(u)a(Y_{\mathrm{src}}(u))b(Y_{\mathrm{tgt}}(T-u))\right| \leq C_{\varepsilon,\chi}L^{-1}\|a\|_\infty\|b\|_\infty \quad (2.11)$$

for $F = F_\lambda, F_0, F_*$. The functions, total and colouring may vary after the threshold. The F_λ part alone means the assertion with that source only.

The precision 1 is sufficient for the counting argument. Section 3 bounds the combined cost of the two inverse occupancy factors and the group normalizations by L^κ, with $\kappa < .975$. The depth of a graph construction used to prove this estimate must be fixed before L, T and the colouring; choosing a separate depth after seeing each colouring would not establish the stated uniformity.

The proposition is not a restatement of the unweighted conclusion. Its weights depend strongly on the prime counts and vanish on missing colours. Sections 3 and 4 prove the passage between the two statements. Its candidate proof from dilation graphs is given in Sections 6 and 7; Proposition 7.3 of *Two smooth summands for every large integer* treats a centred unsigned source at an N^ε-rough total and a terminal function invariant under Q-smooth dilation [7]. Here signed covariance, the primitive mask at a general total, and the dependence on group counts all enter the graph comparison. The reflected Liouville theorem uses only the F_λ instance. The two smoothness centres enter only the joint law and its smoothness variants.

3 Averaging the marking weights

The marked estimate observes integers through every colour. We now show that the average of these observations can be inverted on almost all integers. All results in this section, apart from the explicitly invoked marked estimate, are finite identities or elementary counting estimates.

3.1 Regular colourings have overwhelming probability

Lemma 3.1 (Simultaneous regularity)**.** *Uniformly for T in (2.2), the probability that a random colouring fails (2.10) or has some $V_g \notin [1/2, 3/2]$ is at most*

$$\exp\{-\exp(L^{.27}/3)\} \quad (3.1)$$

for all sufficiently large L.

Proof. Fix a group, an interval inside a dyad $[u, 2u]$, a character and a frequency. The summands

$$(\mathbf{1}_{p\in\mathcal{P}_g} - s_i^{-1})\xi(p)p^{it}$$

are independent centred random variables of modulus at most one. Applying the exponential bounded-summand inequality separately to the real and imaginary parts bounds the probability of exceeding a fixed fraction of the threshold in (2.10) by

$$4\exp\{-cu\exp(-2L^{.27}/100)\}.$$

One may prove this inequality directly from $\mathbb{E}\exp(t(X-\mathbb{E}X)) \le \exp(Ct^2)$ for a real random variable in an interval of bounded length, followed by optimization in t. There are at most $2u+1$ summands and $u \ge \frac{1}{2}e^{L^{.27}}$.

There are $e^{O(L)}$ choices of character, since $\sum_{q\le e^{2L}}\varphi(q) \le e^{4L}$. The effective interval endpoints lie among at most $e^{L^{.46}}$ primes. The number of dyads is $O(L^{.46})$, and the number of groups is $O(\ell)$. For the frequency, take a grid of spacing e^{-5L} in $[-2e^{3L}, 2e^{3L}]$. The derivative of an unweighted sum on $[u,2u]$ has modulus $O(u\log(2u))$, so replacing a frequency by a grid point changes the sum by much less than its threshold. The total number of tests remains $e^{O(L)}$. A union bound is absorbed by the preceding double-exponential estimate, proving (2.10) simultaneously with the probability in (3.1).

Partial summation on each dyad gives, for reciprocal sums, an error $O(\exp(-L^{.27}/100))$ per dyad. Mertens' estimate and

$$\sum_{\substack{p\mid T\\ p>e^{L^{.27}}}}\frac{1}{p} \le \frac{\log T}{\log 2}\,e^{-L^{.27}}$$

then give

$$H_i := \sum_{p\in\mathcal{B}_i(T)} p^{-1} = w_i\ell + o(1), \qquad V_g = H_i/s_i + o(1) = 4/3 + o(1). \tag{3.2}$$

The errors are uniform in T and the group. Thus the mass bounds hold on the same event, after increasing the threshold. □

The random partitions are sampled independently for each proof at a fixed total. No limiting random multiplicative function is constructed. The uniform estimate on every regular colouring allows us to average after the arithmetic estimate; we never multiply an error by the number of colourings.

3.2 The exact occupancy factor

For integers $s \ge 1$ and $k \ge 0$, put

$$q_s(k) = \begin{cases} (k)_s s^{-s}\theta^{k-s}, & k \ge s,\\ 0, & k < s. \end{cases} \tag{3.3}$$

Lemma 3.2 (Averaging the colours of distinct primes)**.** *Assign each of k distinct objects independently to one of s colours, and let K_j be the colour counts. With the zero-factor convention used in (2.5),*

$$\mathbb{E}\prod_{j=1}^{s} K_j\theta^{K_j-1} = q_s(k). \tag{3.4}$$

For $k \ge s$,

$$q_s(k) \ge e^{-s}\theta^{k-s}. \tag{3.5}$$

If two integers are coprime and squarefree on both prime blocks, the expectation of the product of their unnormalized weights is

$$\mathbb{E}_{\mathfrak{c}}\, U_{\mathfrak{c}}(u)U_{\mathfrak{c}}(v) = \prod_{i=1}^{2} q_{s_i}(k_i(u))q_{s_i}(k_i(v)), \tag{3.6}$$

where k_i counts the prime factors in block i.

Proof. If some colour is absent, the product vanishes. Otherwise its power of θ is $k-s$. Expanding $\prod_j K_j$ selects one object of each colour. There are $(k)_s$ ordered selections of distinct objects and each has probability s^{-s} of receiving the prescribed colours. This proves (3.4). Since $(k)_s \geq s!$ and $s! \geq (s/e)^s$, it also proves (3.5).

For coprime u, v, no prime appearing in one integer appears in the other. Their colouring variables are therefore independent. The two prime blocks are also coloured independently. Applying (3.4) at both sites proves (3.6). □

The coprimality condition is essential for this factorization. It will be supplied by $(u, T) = 1$ and $v = T - u$; the noncoprime representations are recovered afterwards.

3.3 Occupancy failures have small ordinary density

Let

$$k_i(v) = \sum_{p \in \mathcal{B}_i(T)} v_p(v), \qquad k(v) = k_1(v) + k_2(v).$$

These block counts are independent of the colouring. Call v typical when

$$k_i(v) \geq s_i \quad (i = 1, 2), \qquad k(v) \leq .18\ell. \tag{3.7}$$

We separately discard integers divisible by the square of a block prime.

Lemma 3.3 (Typical counts)**.** *Uniformly for the allowed T,*

$$\frac{1}{T}\#\{1 \leq v < T : v \text{ is not typical}\} \ll L^{-11/10000}. \tag{3.8}$$

The proportion divisible by a square of a block prime is $O(\exp(-L^{.27}))$.

Proof. The square assertion follows from $\#\{v < T : p^2 \mid v\} \leq T/p^2$ and $\sum_{p > e^{L^{.27}}} p^{-2} \ll e^{-L^{.27}}$. Off this population, multiplicities equal distinct counts. Write ω_i for the latter and $\omega = \omega_1 + \omega_2$.

For the upper tail, set $z_+ = 9/8$ and expand $z_+^{\omega(v)}$ over the squarefree divisors supported on the blocks. All coefficients in this expansion are nonnegative. Counting each set of multiples by $\lfloor T/d \rfloor \leq T/d$ gives

$$\frac{1}{T}\sum_{v<T} z_+^{\omega(v)} \leq \prod_{p \in \mathcal{P}} (1 + (z_+ - 1)/p) \leq \exp((H_1 + H_2)/8) \ll L^{.02}.$$

Consequently

$$\mathbb{P}_{v<T}(\omega > .18\ell) \ll L^{-\delta_+}, \qquad \delta_+ = .18\log(9/8) - .02 > .0012009. \tag{3.9}$$

For a lower tail use $z = 3/4$. The finite identity

$$z^{\omega_i(v)} = \prod_{p \in \mathcal{B}_i(T)} (1 - (1-z)\mathbf{1}_{p|v})$$

has the usual even and odd inclusion–exclusion bounds. Truncate at an even order K of size $C_A\ell$. On averaging over $v < T$, replacing each multiple density by $1/d$ costs at most

$$T^{-1}\exp\{O(KL^{.46})\} = \exp\{-L + o(L)\}.$$

Indeed, the number of choices of at most K primes is bounded by a fixed multiple of $(e^{L^{.46}})^K$. The tail of the corresponding Euler product is bounded by the tail of $\sum_j((1-z)H_i)^j/j!$. Choosing C_A sufficiently large makes it $O_A(L^{-A})$. Thus

$$\frac{1}{T}\sum_{v<T} z^{\omega_i(v)} \le \prod_{p\in\mathcal{B}_i(T)}(1-(1-z)/p) + O_A(L^{-A}) \ll \exp(-(1-z)H_i) + O_A(L^{-A}). \tag{3.10}$$

Since $z^{\omega_i} \ge z^{s_i}$ on $\omega_i < s_i$, this gives

$$\mathbb{P}_{v<T}(\omega_i < s_i) \ll L^{-w_i\gamma+o(1)}, \qquad \gamma = \frac{1}{4} - \frac{3}{4}\log\frac{4}{3}.$$

Here $\gamma > 0.0342$ and $\min_i w_i\gamma > 0.00239$. Both lower tails are smaller than the upper tail in (3.9). Since $\delta_+ > 11/10000$, combining the three tails and the square error proves (3.8). □

These are ordinary density estimates. Restricting to $(v,T) = 1$ only decreases the unnormalized exceptional population. It does not require an independent-prime model for the conditioned integers.

3.4 The inverse has a fixed polynomial cost

Define a colouring-independent correction

$$J_T(v) = \mathbf{1}_{v \text{ typical}} \prod_{i=1}^{2} q_{s_i}(k_i(v))^{-1}, \tag{3.11}$$

with value zero outside the typical set. From (3.5),

$$\|J_T\|_\infty \le \sup_{k\le.18\ell} e^s 256^{k-s} \ll L^{a_J}, \qquad a_J = .12 + .06\log 256 < .453. \tag{3.12}$$

Here $s = .12\ell + O(1)$; retaining the subtraction of s in the exponent is useful. The correction is a function of the block counts and hence is available at either site in (2.4). This bound is independent of the graph depth.

For clarity, use the coarser observations which do not depend on the colouring,

$$\overline{Y}_{\mathrm{src}}(v) = \big((v^\circ)_{\le Q}, k_1(v), k_2(v)\big), \qquad \overline{Y}_{\mathrm{tgt}}(v) = \big(v^\circ, k_1(v), k_2(v)\big).$$

Each is a function of the corresponding observation for every colouring.

Proposition 3.4 (Removal of the weights on primitive pairs)**.** *Assume* (2.11) *for the chosen source F. Uniformly in real tests $|A|, |D| \le 1$ of the two coarser observations,*

$$\frac{1}{T}\left|\sum_{\substack{1\le u<T\\(u,T)=1}} \chi(u/T)F(u)A(\overline{Y}_{\mathrm{src}}(u))D(\overline{Y}_{\mathrm{tgt}}(T-u))\right| \ll_{\varepsilon,\chi} L^{-11/10000}. \tag{3.13}$$

Proof. Insert J_T into each of the two tests in (2.11). Multiply the resulting inequality by $\prod_g V_g^2$ to replace the normalized weights by $U_{\mathfrak{c}}$. The two inverse factors cost $O(L^{2a_J})$. The arithmetic–geometric mean inequality within each block gives the sharper normalization bound

$$\prod_g V_g^2 \le \prod_{i=1}^{2} (H_i/s_i)^{2s_i} \ll L^{.24\log(4/3)}. \tag{3.14}$$

Indeed, $H_i = w_i\ell + o(1)$ and $s_i = 3w_i\ell/4 + O(1)$, uniformly in T. Define

$$\kappa = 2a_J + .24\log(4/3) = .9744649907\ldots < 1. \tag{3.15}$$

The resulting error is $O_{\varepsilon,\chi}(L^{-(1-\kappa)})$, and $1-\kappa > .0255$.

To average over all colourings, reinsert the irregular ones. The unnormalized weights are at most one, so their total contribution in normalized counting is at most

$$O_{\varepsilon,\chi}(L^{2a_J}\mathbb{P}(\mathfrak{c} \text{ irregular})),$$

which is smaller than every inverse power of L by Lemma 3.1. The same bound $O(L^{2a_J})$, together with the square-density estimate, permits deletion of block prime squares at both sites before taking the exact expectation.

On the remaining primitive pairs $u + v = T$, the two integers are coprime. Lemma 3.2 therefore gives the pointwise identity

$$\mathbb{E}_{\mathfrak{c}} U_{\mathfrak{c}}(u)U_{\mathfrak{c}}(v)J_T(u)J_T(v) = \mathbf{1}_{u \text{ typical}}\mathbf{1}_{v \text{ typical}}.$$

The source and the coarser tests are fixed during this expectation. Interchanging a finite sum and expectation removes the weights exactly on the typical set. Finally delete the two typical-set indicators. By Lemma 3.3, their cost is $O_{\varepsilon,\chi}(L^{-11/10000})$ because the source and the remaining tests are bounded. This dominates the preceding errors and proves (3.13). □

The decisive balance is

$$L^{-1}\cdot L^{2a_J}\cdot L^{.24\log(4/3)} = L^{-(1-\kappa)}, \qquad 1-\kappa > .0255.$$

The final, weaker exponent comes from the upper occupancy tail. This balance fixes the required marked precision before any graph depth is chosen.

3.5 How much marked cancellation suffices?

The preceding calculation fixes a power saving for a specified truncation. For qualitative reflected limits one can concentrate the inverse weight near the mean number of prime factors in each block. This gives a weaker sufficient hypothesis on the graph.

Proposition 3.5 (A sufficient exponent for colour averaging)**.** *Keep the two prime blocks and $\theta = 1/256$. For this proposition choose $s_i = \lfloor \alpha w_i \ell \rfloor$ with a fixed $2/3 < \alpha < 1$. Set*

$$K(\alpha) = 2(w_1+w_2)\left\{\alpha + (1-\alpha)\log\frac{1-\alpha}{\theta}\right\}.$$

Suppose that for one of the three sources (2.11) holds with $L^{-\beta}$ in place of L^{-1} and with the same uniformity for these groups. If $\beta > K(\alpha)$, the primitive estimate (3.13) holds with some exponent $c = c(\alpha,\beta) > 0$ in place of $11/10000$. Ordinary cancellation for that source follows after the common-divisor decomposition. At $\alpha = 3/4$ the sufficient condition is $\beta > .24 + .08\log 64$. At $\alpha = 255/256$ it is $\beta > 51/160$.

Proof. For a block, write $H = H_i$, $s = s_i$ and $k = k_i(v)$. The arithmetic–geometric mean bounds the product of its V_g by $(H/s)^s$. Its normalized inverse occupancy cost at one site is thus bounded by

$$\frac{(H/s)^s}{q_s(k)} = \frac{H^s(k-s)!}{k!}\,\theta^{-(k-s)} \qquad (k \ge s).$$

Here $H = w_i\ell + o(1)$ uniformly. Fix a small $\delta > 0$ and retain only $|k/H - 1| \le \delta$ with no block square. Take $\delta < (1-\alpha)/2$. Writing $x = k/H$ and $z = s/H$, Stirling's formula gives the logarithm of the displayed cost divided by H as

$$z + (x-z)\log(x-z) - x\log x - (x-z)\log\theta + O_\alpha(\log H/H).$$

On this band the expression equals $\alpha + (1-\alpha)\log((1-\alpha)/\theta) + O_\alpha(\delta) + o(1)$, since $z \to \alpha$. Apply the bound at the two summands and in both blocks. Given $\eta > 0$, one can choose a fixed δ so their combined cost is at most $L^{K(\alpha)+\eta}$ for large L.

The discarded proportion is $O_\delta(L^{-c_\delta})$ for some $c_\delta > 0$. To see this, let ω_i count distinct block primes. For $v < T$ with normalization $1/T$, the upper-moment expansion in Lemma 3.3 bounds the mean of $e^{h\omega_i}$ by $\exp(H(e^h - 1))$ for $h > 0$. For $h < 0$ its even inclusion–exclusion truncation there gives the same upper bound plus $O_A(L^{-A})$ for each fixed h and A. Apply Markov with $h = \log(1+\delta)$ and $h = \log(1-\delta)$ on the respective tails. The two exponents are $-H[(1\pm\delta)\log(1\pm\delta) \mp \delta]$, both strictly negative. Choose A after δ to absorb the lower-tail error. Block squares have exponentially small density.

The new counts also lie in the ranges of (6.1) for all sufficiently large L, since $s_i/\ell \to \alpha w_i$. The proof of Lemma 3.1 still applies. Each group has mass $1/\alpha + o(1)$ on the same concentration event, strictly between $1/2$ and $3/2$. Use the two inverse factors on the retained bands as in Proposition 3.4. Their sizes are polynomial in L, so irregular colourings and square divisors still have negligible cost. Take $0 < \eta < \beta - K(\alpha)$. The marked error is then $O(L^{-(\beta-K(\alpha)-\eta)})$ and the ordinary exceptional set costs $O_\delta(L^{-c_\delta})$. Both absorb the $O(\log L)$ divisor sum in Lemma 4.1. Finally, the derivative of $K(\alpha)/(2(w_1+w_2))$ is $-\log((1-\alpha)/\theta)$, so its minimum on the stated interval occurs at $\alpha = 1-\theta$ and equals $1-\theta$. Since $w_1 + w_2 = .16$, this gives $51/160$. □

The explicit rate in Proposition 3.4 continues to use its original choice of colours. Proposition 3.5 measures the loss in a different truncation and supplies a sufficient marked estimate. It gives neither a necessary exponent for reflected cancellation nor an estimate on a graph operator.

4 Recovering ordinary reflected averages

The primitive-pair estimate has the right normalization for summing over common divisors. We first give the general local version, which retains the smoothness cutoff N^ε and the small-prime factors needed for the extensions.

Lemma 4.1 (Common-divisor tails)**.** *For $L = \log N$ and $G = \exp(L^{.20})$,*

$$\frac{1}{N}\#\{1 \le n < N : (n,N) > G\} \ll_A L^{-A} \quad \textit{for every fixed } A > 0. \tag{4.1}$$

Also $\sum_{d|N} d^{-1} \ll \log L$.

Proof. For $t = 1/\log L \le 1/2$, as holds for large L, counting multiples gives

$$\frac{1}{N}\#\{n < N : (n,N) > G\} \le \sum_{\substack{d|N \\ d>G}} d^{-1} \le G^{-t}\sum_{d|N} d^{-1+t}.$$

The last divisor sum is at most $\prod_{p|N}(1-p^{-1+t})^{-1}$. For $p \le L^2$ we have $p^t \le e^2$, and the logarithm of this part of the product is $O(\log\log L)$ by Mertens. There are at most $L/\log 2$ larger prime divisors, and each contributes $O(L^{-2+2t})$ to the logarithm. Hence

$$\sum_{d|N} d^{-1+t} \ll (\log L)^C$$

for an absolute C. Multiplication by $G^{-t} = \exp(-L^{.20}/\log L)$ proves (4.1). For the final assertion use $t = 0$ and the sharper Mertens product at L^2; the large-prime part is bounded. □

Proposition 4.2 (Local estimates at the original total)**.** *Assume* (2.11) *for the source corresponding to H below. Put $S(v) = \mathbf{1}_{P^+(v)\le N^\varepsilon}$ and $r = \rho(1/\varepsilon)$. For fixed real $\chi \in C_c^\infty((0,1))$, uniformly in real functions $|a|, |b| \le 1$ on the positive integers,*

$$\frac{1}{N}\left|\sum_{1\le n<N} \chi(n/N)H(n)a(n_{\le B})b((N-n)_{\le B})V(N-n)\right| \ll_{\varepsilon,\chi} L^{-1/1000}, \quad (4.2)$$

where $H = \lambda, S-r, S\lambda$ respectively uses F_λ, F_0, F_*, and V may be any of $1, S, \lambda, S\lambda$.

Proof. First keep the centred source F of (2.8) in place of H. Partition n by $d = (n, N)$, and write $n = du$, $N = dT$, $N - n = dv$ with $(u, T) = 1$, $v = T - u$. By Lemma 4.1, the terms with $d > G$ have negligible ordinary density. For $d \le G$, (2.2) holds and every prime factor of d is below B.

The small-prime masks become

$$a_d(u_{\le B}), \qquad b_d(v_{\le B}), \qquad a_d(w) = a(dw), \quad b_d(w) = b(dw),$$

still bounded functions of their separate local arguments. Moreover, $S(du) = S(u)$ and $S(dv) = S(v)$ with the *unchanged* cutoff $y = N^\varepsilon$. The factors $\lambda(d)$ in the signed sources and in $V(dv)$ are constant on the stratum. By (2.9), they do not affect absolute estimates.

Each source mask is measurable in $\overline{Y}_{\mathrm{src}}$ and each target mask multiplied by V is measurable in $\overline{Y}_{\mathrm{tgt}}$. Proposition 3.4 therefore bounds the stratum by $O_{\varepsilon,\chi}(TL^{-11/10000})$. Since $T = N/d$, their sum divided by N is at most

$$O_{\varepsilon,\chi}\left(L^{-11/10000}\sum_{d|N} d^{-1}\right) \ll_{\varepsilon,\chi} L^{-11/10000}\log L.$$

This is $O_{\varepsilon,\chi}(L^{-1/1000})$.

It remains to replace the centred sources by the displayed H. No replacement is needed for F_λ. For F_*, the pointwise error is at most $\|\mathcal{M}_-\|_\infty \ll_\varepsilon L^{-.49}$. For F_0, discard the integers with $n^\flat > D_* := \exp(L^{.49}\ell^2)$. Lemma 5.1 shows that their ordinary density is smaller than every inverse power of L. On the support of χ, the remaining integers satisfy

$$\log(n^\sharp) = L + O_\chi(L^{.49}\ell^2),$$

so (2.6) and the local Lipschitz continuity of ρ give

$$\mathcal{M}_0(n^\sharp) = r + O_{\varepsilon,\chi}(L^{-.02}\ell^2).$$

Both centre errors are smaller than $L^{-1/1000}$. This proves (4.2), uniformly in the local masks. □

Lemma 4.3 (Uniform absolute errors over intervals)**.** *Let $h_{N,\alpha}(n)$, $1 \le n < N$, be a family of complex arrays with $|h_{N,\alpha}(n)| \le C$, and let $|c_{N,\alpha}| \le C$. Suppose that, for every fixed $\chi \in C_c^\infty((0,1))$,*

$$\sup_\alpha \left| \frac{1}{N} \sum_{1 \le n < N} \chi(n/N) h_{N,\alpha}(n) - c_{N,\alpha} \int_0^1 \chi(t)\, dt \right| \longrightarrow 0.$$

Then

$$\sup_{\alpha,\ I \subseteq (0,1)} \left| \frac{1}{N} \sum_{n/N \in I} h_{N,\alpha}(n) - c_{N,\alpha} |I| \right| \longrightarrow 0, \tag{4.3}$$

where the supremum is over intervals, including $(0,1)$. The constants $c_{N,\alpha}$ need not converge.

Proof. For a fixed interval, approximate its indicator by a smooth function which agrees with it outside boundary intervals of arbitrarily small total length. The bound on h controls the error by C times this length, plus $O_C(N^{-1})$; the bound on c controls the corresponding integral error. Endpoints at 0 or 1 are handled by trimming boundary intervals first. Thus the asserted convergence holds uniformly in α for every fixed interval.

Now fix an integer $K \ge 1$. There are finitely many intervals with endpoints in $\{0, 1/K, \dots, 1\}$, so the maximum of their discrepancies tends to zero uniformly in α. Any interval I differs from one of these grid intervals in a set of length at most $2/K$, apart from its endpoints. Therefore its discrepancy is at most the grid maximum plus $4C/K + O_C(N^{-1})$. First let $N \to \infty$ with K fixed, and then let $K \to \infty$. This proves (4.3). □

The normalization in this lemma is by N. It allows an interval $I = I_N$ to move or shrink, but it supplies no relative error $o(N|I_N|)$ when $|I_N| \to 0$.

Proof of Theorem 1.1. Take $H = \lambda$, $V = \lambda$, and $a = b = 1$ in Proposition 4.2. It gives the cancellation with any fixed smooth interior weight. Choose $0 \le \chi_\delta \le 1$ equal to one on $[\delta, 1-\delta]$ and supported in $(0,1)$. The unweighted complement has at most $2\delta N + 2$ terms. Letting $N \to \infty$ first and then $\delta \downarrow 0$ proves (1.3).

The two one-site Liouville sums with a fixed smooth weight are $o(N)$ by the prime number theorem and partial summation, including after reflection. Expanding the two sign indicators therefore gives their smooth weighted mean $\frac{1}{4} \int \chi + o(1)$. Apply Lemma 4.3 to their product, with $c = 1/4$. It proves (1.4) uniformly over all intervals. Only the F_λ part of Proposition 2.2 has been used in this proof. □

Remark 4.4 (An alternative lift for the unmodulated pair). For the product $\lambda(n)\lambda(N-n)$ alone, one can avoid the quantitative divisor tail. Normalize the primitive estimate by $\varphi(T)$, using $T/\varphi(T) \ll \log\log T$. Then sum over $T \mid N$ and use $\sum_{T|N} \varphi(T) = N$. Primitive totals $T > M$ have uniformly small normalized error, while $T \le M$ contribute $O(M^2/N)$. Taking $N \to \infty$ and then $M \to \infty$ gives the result. The common-scale argument above is used because it also preserves smoothness and the local multipliers.

The limits just used do not turn the smooth-weight estimate into an explicit logarithmic error for a sharp full interval. No such global error is asserted here. They do give the ordinary $o(N)$ conclusion at every total.

5 One-variable cancellation and the two centres

The graph comparison will ask for cancellation against an arbitrary bounded function of the Q-smooth part. The prime number theorem does not give that assertion by a direct partial summation in the original variable: the coefficient need not have any regularity. The appropriate order is to fix the small factor first, and then sum over the rough factor. We give that argument for all three sources, including the signed smoothness source.

Throughout this section the parameters are those of (2.1), even when a test is evaluated at a different physical scale. Set

$$h_Q(v) = \mathbf{1}_{P^-(v)>Q}, \qquad D_Q = \prod_{p\le Q}(1-p^{-1}), \qquad D_* = \exp(L^{.49}\ell^2),$$

with $h_Q(1) = 1$ and $P^+(1) = 1$. In particular, $D_Q^{-1} \asymp L^{.49}$.

5.1 Rough factors and smooth reciprocal tails

Lemma 5.1. *For every fixed $A > 0$,*

$$\sum_{\substack{a>D_*\\ P^+(a)\le Q}} \frac{1}{a} \ll_A L^{-A}. \tag{5.1}$$

Writing $R_Q(x) = \sum_{v\le x} h_Q(v)$, we have

$$|R_Q(x) - D_Q x| \le \lceil D_*\rceil + 1 + x \sum_{\substack{a>D_*\\ P^+(a)\le Q}} \frac{1}{a}. \tag{5.2}$$

Furthermore, for $1 \le X \le e^{2L}$,

$$\sum_{v\le X} \frac{h_Q(v)}{v} = D_Q \log X + O(L^{.49}\ell^2). \tag{5.3}$$

The proportion of integers $n \le X$ with $n^\flat > D_*$ is $O_A(L^{-A})$, uniformly in $X \ge 1$.

Proof. Take $\delta = 1/\log Q$. Rankin's inequality bounds the left side of (5.1) by

$$D_*^{-\delta} \prod_{p\le Q}(1-p^{-1+\delta})^{-1}.$$

The product is $O(\log Q)$. Indeed, its logarithm differs from that of the ordinary Mertens product by $O(1)$: the first-order difference is at most $e\delta\sum_{p\le Q}(\log p)/p = O(1)$, and the terms of order at least two form a uniformly convergent series. Since $D_*^{-\delta} = e^{-\ell^2}$, the tail has any fixed logarithmic saving.

Use the exact sieve identity

$$h_Q(v) = \sum_{a\mid v} \mu(a)\mathbf{1}_{P^+(a)\le Q}.$$

In the count of multiples, the rounding errors for $a \le D_*$ sum to at most $\lceil D_*\rceil$; the remaining errors are bounded by x/a. Passing from an integer endpoint to x costs at most one. This proves (5.2). Partial summation then proves (5.3): before D_* use the elementary bound $|R_Q(t) - D_Q t| \le t$, and afterwards use (5.2) with arbitrarily high tail precision. Finally, any integer with $n^\flat > D_*$ is divisible by a Q-smooth integer $a > D_*$. Counting its multiples by X/a and using (5.1) proves the density assertion. □

These elementary estimates also imply the following comparison. If v is a right-continuous function of bounded variation supported in $(0, X)$, then

$$\left|\sum_n h_Q(n)v(n) - D_Q\int v(x)\,dx\right| \le \Delta_Q(X)(2\|v\|_\infty + \operatorname{Var} v), \tag{5.4}$$

where $\Delta_Q(X) = \sup_{0\le t\le X}|R_Q(t) - D_Q t|$. To see the convention at a jump, let $E(t) = R_Q(t) - D_Q t$ on $[0, X]$. The integration-by-parts formula uses $v(t)\,dE(t)$ and $E(t-)\,dv(t)$. Thus the sum in (5.4) has the prescribed right-hand values. For $X \ge e^{cL}$, with $c > 0$ fixed, (5.2) gives $\Delta_Q(X) \ll_{A,c} XL^{-A}$.

5.2 The unsigned and signed kernels

Put $\ell_y = \log y$ and let

$$h(t) = \frac{\mathbf{1}_{[1,\infty)}(t)}{t}, \qquad A_j = h^{*j} \quad (j \ge 1).$$

The convolution is additive, and $A_1(1) = 1$. Define

$$I_j(u) = \int_{\substack{t_1,\dots,t_j\ge 1\\ t_1+\cdots+t_j\le u}} \frac{dt_1\cdots dt_j}{t_1\cdots t_j}, \qquad I_0(u) = \mathbf{1}_{[0,\infty)}(u).$$

The identities

$$I_j(u) = \int_{-\infty}^{u} A_j(t)\,dt, \qquad uA_j(u) = jI_{j-1}(u-1) \quad (u > 0) \tag{5.5}$$

follow by symmetry in the j simplex coordinates. They show that each A_j vanishes below j and has bounded supremum and variation on any fixed bounded upper range.

For $e^{L/2} \le x \le e^{2L}$, define

$$\mathcal{M}_0(x) = 1 + \frac{1}{D_Q\ell_y}\sum_{j=1}^{K}\frac{(-1)^j}{j!}\sum_{v\ge 1}\frac{h_Q(v)}{v}A_j\left(\frac{\log(x/v)}{\ell_y}\right), \tag{5.6}$$

$$\mathcal{M}_-(x) = \frac{1}{D_Q\ell_y}\sum_{j=1}^{K}\frac{1}{j!}\sum_{v\ge 1}\frac{a_Q(v)}{v}A_j\left(\frac{\log(x/v)}{\ell_y}\right), \qquad a_Q(v) = h_Q(v)\lambda(v), \tag{5.7}$$

where K is any fixed integer greater than $2/\varepsilon$. The extra terms are zero, so the definitions do not depend on this choice. Both sums are locally finite. Outside the displayed x-interval extend $\mathcal{M}_0$ by $r = \rho(1/\varepsilon)$ and $\mathcal{M}_-$ by zero.

The unsigned model is the invariant centre of [7, Section 4], specialized to $\log Q = L^{.49}$. The signed model has two differences: it has no constant term, and its displayed coefficients are all positive. The first comes from rough Liouville cancellation; the second from the compensation between the inclusion–exclusion sign and the signs of the removed primes.

Lemma 5.2 (Size and variation of the centres)**.** *Uniformly on $e^{L/2} \le x \le e^{2L}$,*

$$\mathcal{M}_0(x) = \rho\left(\frac{\log x}{\log y}\right) + O_\varepsilon(L^{-.02}\ell^2), \qquad \|\mathcal{M}_-\|_\infty \ll_\varepsilon L^{-.49}. \tag{5.8}$$

Both models have total variation $O_\varepsilon(1)$ on this interval and remain bounded after extension. For large N, $0 \le \mathcal{M}_0 \le 1$.

Proof. Apply (5.3) to the bounded-variation kernel $A_j(u - \log v/\ell_y)$. It gives

$$\sum_{v\ge 1} \frac{h_Q(v)}{v} A_j\left(u - \frac{\log v}{\ell_y}\right) = D_Q \ell_y I_j(u) + O_\varepsilon(L^{.49}\ell^2).$$

For $j = 1$, the summation ends at $v = x/y$ with that endpoint included; ordinary partial summation with the increasing function $(u - \log t/\ell_y)^{-1}$ gives the same bound while retaining its right-hand endpoint value. Substitution into (5.6) gives the finite simplex sum $\sum_{j\ge 0}(-1)^j I_j(u)/j!$ with error $O_\varepsilon(L^{-.02}\ell^2)$. By (5.5), this sum is one on $[0,1]$ and satisfies $u\rho'(u) = -\rho(u-1)$ for $u > 1$; it is therefore the Dickman function. Its values on the fixed range $[1/(2\varepsilon), 2/\varepsilon]$ are bounded away from both zero and one. This proves the unsigned assertions about values.

For the signed model, the exact Dirichlet convolution identity is

$$a_Q = \lambda * b_Q, \qquad b_Q(d) = \mu^2(d)\mathbf{1}_{P^+(d)\le Q}. \tag{5.9}$$

At a prime $p \le Q$, the coefficients of every positive power cancel as $\lambda(p^j) + \lambda(p^{j-1}) = 0$; at $p > Q$, the coefficient remains $\lambda(p^j)$. Let $U_0 = e^{L^{.51}}$. Truncate d in (5.9) at D_*. The tail is at most t times (5.1). For $d \le D_*$ and $U_0 \le t \le e^{2L}$, eventually $\log(t/d) \ge \frac{1}{2}L^{.51}$. The classical prime number theorem for λ, with an arbitrarily high logarithmic precision, and $\sum b_Q(d)/d \ll L^{.49}$ give

$$\sum_{v\le t} a_Q(v) \ll_A tL^{-A} \qquad (U_0 \le t \le e^{2L}). \tag{5.10}$$

The required classical estimate follows, for example, from the prime number theorem with its usual zero-free-region remainder for μ and the identity $\lambda = \mu * \mathbf{1}_{\text{squares}}$; see [10].

Below U_0, the rough reciprocal mass is bounded by $\prod_{Q<p\le U_0}(1-p^{-1})^{-1} \ll L^{.02}$. Above U_0, partial summation of (5.10) then proves

$$\sup_{t\le e^{2L}} \left|\sum_{v\le t} \frac{a_Q(v)}{v}\right| \ll L^{.02}, \qquad \sum_{v\le e^{2L}} \frac{|a_Q(v)|}{v} \ll L^{.51}. \tag{5.11}$$

As $D_Q\ell_y \asymp_\varepsilon L^{.51}$, integration against the first bound and the bounded-variation kernels in (5.7) gives $\|\mathcal{M}_-\|_\infty \ll_\varepsilon L^{-.49}$.

For the variation of either centre, sum the variations of its finitely many kernels against the absolute reciprocal coefficients. The second bound in (5.11) also bounds the unsigned coefficient mass, so division by $D_Q\ell_y$ leaves $O_\varepsilon(1)$. The bounded extensions add only bounded endpoint jumps. □

5.3 Cancellation with arbitrary small-factor coefficients

Proposition 5.3 (The principal-character input)**.** *Let $A > 0$, $B_1 \ge 0$ and $C \ge 1$ be fixed. Suppose that*

$$e^{3L/4} \le Z \le e^{5L/4}, \quad \operatorname{supp} w \subset [Z/C, CZ], \quad \|w\|_\infty + \int |w'(x)|\,dx \le L^{B_1}, \qquad |c(a)| \le L^{B_1}.$$

Here $w \in C_c^1(\mathbb{R})$ may be complex, and c is any function on the Q-smooth integers. For each $F \in \{F_\lambda, F_0, F_*\}$,

$$\sum_{n\ge 1} F(n)c(n^\flat)w(n) \ll_{A,B_1,C,\varepsilon} ZL^{-A}. \tag{5.12}$$

The functions c, w may depend on all auxiliary parameters subject to these bounds.

Proof. We first prove the rough-factor versions with $e^{5L/8} \le Z \le e^{11L/8}$. For F_λ this follows from (5.10) and partial summation. For the other sources, expand $S(v)$ by inclusion–exclusion over its distinct prime divisors exceeding y. The number of selected primes is at most the fixed integer K on the support of w. In the unsigned case the coefficient is $(-1)^j/j!$ and the residual factor is $h_Q(v)$. In the signed case

$$h_Q(vp_1\cdots p_j)\lambda(vp_1\cdots p_j) = (-1)^j a_Q(v)$$

even if a selected prime divides the residual integer. Thus the coefficient becomes $1/j!$, as in (5.7).

Allowing repeated prime coordinates costs at most

$$O_{\varepsilon,C}\Big(ZL^{B_1+K+1}/y\Big),$$

by summing p^{-2} over a repeated pair and reciprocal masses over the other coordinates. This error has any fixed logarithmic saving. Replace each remaining prime measure by $dt/\log t$, one coordinate at a time. When the other coordinates and the residual integer are fixed, write their product as U. The test of the selected coordinate is $w(Ut)$ on $y < t \ll Z/U$, with the same supremum and integral of its derivative as w. The classical prime number theorem and partial summation bound its replacement error by

$$O_C\left(\frac{Z}{U}L^{B_1+1}e^{-c_\varepsilon\sqrt{L}}\right).$$

Summing reciprocal masses of the remaining coordinates costs only a fixed power of L. This justifies the replacement uniformly also in the signed case, because $|a_Q(v)| \le h_Q(v) \le 1$ and the residual factor is fixed before each replacement.

For $j \ge 1$, the substitutions $t_i = e^{\ell_y u_i}$ and then $x = ve^{\ell_y(u_1+\cdots+u_j)}$ give the density identity

$$\int_{t_1,\dots,t_j>y} w(vt_1\cdots t_j)\prod_{i=1}^{j}\frac{dt_i}{\log t_i} = \frac{1}{v\ell_y}\int A_j\left(\frac{\log(x/v)}{\ell_y}\right)w(x)\,dx. \tag{5.13}$$

The term with $j=0$ is $D_Q\int w + O_A(ZL^{-A})$ in the unsigned case, by (5.4), and $O_A(ZL^{-A})$ in the signed case, by (5.10). Consequently the two uncentred rough sums equal $D_Q\int \mathcal{M}_0 w$ and $D_Q\int \mathcal{M}_- w$, respectively, up to $O_A(ZL^{-A})$. By Lemma 5.2 and (5.4), these integrals are the corresponding rough sums of the centres with the same precision. Subtraction proves

$$\sum_v h_Q(v)F(v)w(v) \ll_{A,B_1,C,\varepsilon} ZL^{-A}$$

for all three sources in the inner scale range. This argument also explains the endpoint convention in the kernel definitions. It is the signed extension of the unsigned centering argument in [7, Section 4].

Finally write $n = av$ with $a = n^\flat$ and $h_Q(v) = 1$. The source becomes $F(v)$ in the unsigned case and $\lambda(a)F(v)$ in either signed case. The factor $\lambda(a)$ can be absorbed into $c(a)$. Cut a at D_*. The tail contributes at most $O_C(ZL^{2B_1}\sum_{a>D_*,\,P^+(a)\le Q} a^{-1})$. For $a \le D_*$, the test $w(av)$ has scale Z/a in the inner range, the same variation bound, and the same fixed outer parameters L, Q, y and centres. Apply the rough estimate with precision chosen after A, B_1 and sum a^{-1}; its total mass is $D_Q^{-1} \ll L^{.49}$. The resulting bound is (5.12). □

This proposition licenses arbitrary coefficients of $n^\flat$ only after the selected graph prime has been extracted. At the original source, the observation must still be restricted as in (2.4). Section 7.1, Lemma 7.2, checks that extraction explicitly.

5.4 Signs and smoothness at one summand

For $0 \le t \le 1/4$ let $S_t(n) = \mathbf{1}_{P^+(n) \le N^t}$, put $r(0) = 0$ and $r(t) = \rho(1/t)$ for $t > 0$, and set $\pi_t(1) = r(t)$, $\pi_t(0) = 1 - r(t)$. Hildebrand studied the distribution of Ω on smooth integers [26]; de la Bretèche and Tenenbaum treat friable sums of oscillating arithmetic functions [27]. The local weight here can be any bounded function of $n_{\le B}$ and can vary with N. The proof below derives that uniformity directly from the preceding one-variable estimate.

Proposition 5.4 (One-site signs with a local weight)**.** *With $B = \exp((\log N)^{.26})$, one has*

$$\sup_{t,I,i,\sigma,a} \frac{1}{N} \left| \sum_{\substack{1 \le n < N \\ n/N \in I}} a(n_{\le B}) \left(\mathbf{1}_{S_t(n)=i,\, \lambda(n)=\sigma} - \frac{\pi_t(i)}{2} \right) \right| \longrightarrow 0.$$

The supremum ranges over $0 \le t \le 1/4$, intervals $I \subseteq (0,1)$, $i \in \{0,1\}$, $\sigma \in \{-1,1\}$ and all real functions $|a| \le 1$. The proof uses the classical prime number theorem.

Proof. Fix $0 < t < 1/4$ and $\chi \in C_c^\infty((0,1))$. Use Proposition 5.3 with $y = N^t$, $Z = N$, $w(x) = \chi(x/N)$ and $c(v) = a(v_{\le B})$. Since $B < Q$, one has $c(n^\flat) = a(n_{\le B})$. The three centred sums divided by N tend to zero uniformly in a. The bound for $\mathcal{M}_-$ in (5.8) allows F_* to be replaced by $S_t\lambda$. For n in the support of the weight with $n^\flat \le D_*$, we have $\log(n^\sharp) = L + O_\chi(L^{.49}\ell^2)$. The estimate for $\mathcal{M}_0$ therefore allows F_0 to be replaced by $S_t - r(t)$. Lemma 5.1 bounds the discarded proportion. Expanding the indicator as

$$\mathbf{1}_{S_t=i,\,\lambda=\sigma} = \tfrac{1}{2}(iS_t + (1-i)(1-S_t))(1+\sigma\lambda)$$

gives the asserted centring with the smooth weight. Apply Lemma 4.3 with constant zero to the difference in parentheses multiplied by a. This proves uniformity in I for each fixed interior t.

Write $D_N(t) = N^{-1}\#\{1 \le n < N : S_t(n) = 1\}$. Taking $a = 1$ and summing the signs gives $D_N(t) \to r(t)$ in the interior. At zero $D_N(0) = 1/N$. For $0 < u < b = 1/4$, counting multiples of primes gives

$$0 \le D_N(b) - D_N(u) \le \sum_{N^u < p \le N^b} \frac{1}{p} = \log(b/u) + o_{u,b}(1).$$

Thus $D_N(b) \to r(b)$ by continuity. Also $r(t) \to 0$ as $t \downarrow 0$. Choose a fixed finite partition of $[0,b]$ with at least one interior endpoint and oscillation of r at most η on each cell. Move each t to an interior endpoint of its cell. The change of the empirical indicator, tested on any interval with any allowed a, is bounded after normalization by the unweighted mass of the marginal band, at most $\eta + o(1)$. Its centring changes by at most $\eta/2$. The finitely many interior estimates hold simultaneously. Let $N \to \infty$ and then $\eta \downarrow 0$. □

In particular the signs are equidistributed among the N^t-smooth integers for each fixed $t > 0$ in this range. At two sites even a target multiplier $\lambda(N-n)$ depends on factors above B. Proposition 5.4 gives no estimate against that multiplier.

6 A parametric transfer from the ideal graph

This section isolates the graph statement needed for the marked estimate. It is a statement about finite weighted operators: no Liouville correlation occurs in its hypothesis. We give

the extension of the physical-moment argument of Howe, Huffman, Kohlmeyer and Kruer [7, Sections 5–7] that is required when a group of primes is not an interval. The proof below uses their finite memory identities and quantitative estimates, and records the hypotheses on which those estimates depend. This is a mathematical adaptation of their construction; their existing formalization is not a formalization of this adaptation.

6.1 Finite prime families and the two operators

Let L tend to infinity and put $T_0 = \log L$ and $\vartheta = 1/256$. Let $\mathcal{G} = \mathcal{S} \sqcup \mathcal{B}$ be a finite index set and let $\mathcal{P}_g$ be pairwise disjoint finite sets of primes. Write $s_S = |\mathcal{S}|$, $s_B = |\mathcal{B}|$, $s = s_S + s_B$, and

$$V_g = \sum_{p \in \mathcal{P}_g} \frac{1}{p}, \qquad \mu_g(p) = \frac{1}{V_g p}.$$

The following data are imposed uniformly:

$$\begin{aligned}
&\exp(L^{.27}) < p \le \exp(L^{.36}) && (p \in \mathcal{P}_g,\ g \in \mathcal{S}),\\
&\exp(L^{.39}) < p \le \exp(L^{.46}) && (p \in \mathcal{P}_g,\ g \in \mathcal{B}),\\
&\tfrac{1}{2} \le V_g \le \tfrac{3}{2}, \qquad .045T_0 \le s_S \le .09T_0, \qquad .035T_0 \le s_B \le .07T_0. && (6.1)
\end{aligned}$$

There is no ordering or separation condition between two groups in the same block. Their prime sets, rather than their interval envelopes, must be disjoint.

Fix $J \ge 1$ and put $M = J + 1$. A pattern ν specifies subsets $C_{\nu,g} \subseteq \{1, \dots, J\}$. Put

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

and require $C_{\nu,g} = \varnothing$ for $g \in \mathcal{S}$. For source labels $\mathbf{p}$ and free labels $\mathbf{r}$ define

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

Free labels are drawn with replacement. Coefficients $c_\nu(\mathbf{p}_B, \mathbf{p}'_B, \mathbf{r}_B)$ depend only on the indicated high-block labels.

The ideal operator $\mathcal{T}_\zeta(\Theta)$ acts on

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

In the row for $\mathbf{p}_B$, copy the shared coordinates to the target, integrate the unshared target coordinates and the free coordinates with their product laws, and multiply the integrand by

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

Sum over ν and apply at each end the averaging projection over $\prod_{g \in \mathcal{B}} \mathfrak{S}_M$. The two endpoint averages are independent, and their permutations are independent between groups. Here and below $e(t) = \exp(2\pi i t)$. This specifies the row normalization of the ideal operator completely.

For a finite integer window $\mathcal{W}$, the physical vertices are

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

Give each vertex mass $\tau = \prod_g V_g^{-M}$, and set

$$d_\vartheta(z, \mathbf{p}) = \vartheta^{(\omega_\mathcal{P}(z) - Ms)/2}, \qquad \omega_\mathcal{P}(z) = \sum_g \sum_{p \in \mathcal{P}_g} \mathbf{1}_{p|z}.$$

The raw physical kernel for a pattern is

$$K_{\nu,k}^{\zeta,\theta}(x,y) = d_\vartheta(x)d_\vartheta(y)\prod_g V_g^{-t_{\nu,g}}\sum_{\mathbf{r}}\left(\prod_g\prod_{i\in C_{\nu,g}}\mu_g(r_{g,i})\right) \times c_\nu(\mathbf{p}_B,\mathbf{p}'_B,\mathbf{r}_B)D_\nu^{i\zeta}e(\theta k D_\nu)\,\mathcal{I}_\nu(x,y,\mathbf{r}), \tag{6.2}$$

where $\mathcal{I}_\nu$ requires $z' = z + kD_\nu$, agreement of the shared labels, and that each free label differ from every source and target label in its group. Sum these kernels and average independently at the two endpoints over the full product $\prod_{g\in\mathcal{G}}\mathfrak{S}_M$ to obtain $\mathcal{A}_k^{\zeta,\theta}$. The row action is the ordinary sum $Af(x) = \sum_y K(x,y)f(y)$; its adjoint is taken in $L^2(\mathcal{V}_\mathcal{W},\tau)$. Let u_z be the indicator of the fibre of position z. In Euclidean coordinates it has value $\sqrt{\tau}$ on that fibre, rather than value one.

6.2 The precise transfer statement

Proposition 6.1 (Transfer for finite prime families, candidate graph argument)**.** *The following implication* G *holds for the data above. Given* $A > 0$*, an integer* $m \geq 1$*, and* $C_1 \geq 0$*, there is* $J_0 \geq m$ *such that, for each fixed* $J \geq J_0$*, a threshold in* L *can be chosen before the prime family and the pattern coefficients. For every family satisfying* (6.1)*, and every finite pattern family satisfying*

$$|C_{\nu,g}| \leq m \quad (g\in\mathcal{B}), \qquad \sum_\nu \|c_\nu\|_\infty \leq L^{C_1},$$

the implication

$$\sup_{\Theta\in\mathbb{R}}\|\mathcal{T}_\zeta(\Theta)\| \leq L^{-(A+13)} \quad\Longrightarrow\quad \sum_{z\in[X,2X)\cap\mathbb{Z}}\langle u_z,(\mathcal{A}_k^{\zeta,\theta}(\mathcal{A}_k^{\zeta,\theta})^*)^R u_z\rangle \leq XL^{-2AR} \tag{6.3}$$

holds for

$$R = \lceil L^{.52}\rceil, \qquad |\zeta| \leq L^2, \qquad \theta\in\mathbb{R}, \qquad |k| \leq \lfloor e^L\rfloor,$$
$$p\nmid k \quad (p\in\bigcup_g\mathcal{P}_g), \qquad e^{.75L} \leq X \leq e^{1.25L}, \quad X\in\mathbb{N}.$$

For this assertion take

$$Q_S = \prod_{g\in\mathcal{S}}\prod_{p\in\mathcal{P}_g} p, \quad U = \lfloor e^L\rfloor, \quad H = 2RU^{2Ms+1}, \quad \mathcal{W} = [-H,-H+2Q_S)\cap\mathbb{Z}.$$

The coefficients are fixed before ζ,θ,k,X*. The implication may be applied simultaneously in these variables when its ideal-norm premise holds simultaneously. In applications we choose one coefficient family before either the source or the target reader, and use that family for all the readers.*

The position window contains the source fibres and the negative target window used below. Indeed $H = \exp((2Ms+1)L+o(L))$*, whereas* $X \leq e^{1.25L}$*. Thus, for each fixed* $C > 0$*,* $CX < H$ *for all sufficiently large* L*. Together with* $2H \leq Q_S$ *and* $2X \leq Q_S$*, this puts both* $[X,2X)$ *and* $[-CX,-1]$ *inside* $\mathcal{W}$*. An alternating path of* $2R$ *edges starting in* $[X,2X)$ *stays in* $\mathcal{W}$ *because each displacement is at most* U^{2Ms+1}*. The initial and terminal label arrays in the root moment are summed independently; the root is the whole position fibre, not a single labelled vertex.*

6.3 Proof of the physical transfer

We now prove the physical part of Proposition 6.1. The proof explains why the low-rank argument keeps the signed sum of patterns, and why the errors in passing to residues and

truncating memory remain small after the absolute weights of all paths have been summed. The finite identities used below are those in [7, Section 5]. We state their current alphabets and normalizations explicitly. No assertion about an interval of primes is used in this proof.

The table indicates where the finite identities from the cited construction enter. The bounds for the dispersed groups are derived below. In particular the signed sum of all unchanged edges is needed at low rank.

HHKK formulas	Finite identity or bound	Check for the present groups
(5.13)–(5.22)	Tilted masses and factorial adjoint	Tilt only high primes; the factorial mass is unchanged.
(5.23)–(5.34)	Store–translate–retrieve edge	Full free laws and independent products of permutations at both ends.
(5.35)–(5.53)	Completion and small complete-edge norm	Low omissions have M^{s_S} choices and are distinct below P_*.
(5.55)–(5.57)	Boundary normalization	Chinese remainder identity for the current disjoint prime sets.
(5.58)–(5.78)	High- and low-rank bounds	Prescribed fresh atoms; phase Schur bounds; signed unchanged edges.
(5.83)	Physical-to-residue error	Relative Bonferroni error and total rounding cost after the path sum.
(5.84)	Memory cutoff error	Positive birth-count estimate after summing all histories.

6.3.1 Alphabets, symmetrization, and memory masses

Retain the disjoint prime sets $\mathcal{P}_g$, the partition $\mathcal{G} = \mathcal{S} \sqcup \mathcal{B}$, and the bounds (6.1). Thus $s_S \asymp \log L$, $s_B \asymp \log L$, $1/2 \le V_g \le 3/2$, low primes lie between $e^{L^{.27}}$ and $e^{L^{.36}}$, and high primes lie between $P_* = e^{L^{.39}}$ and $e^{L^{.46}}$. Fix the replacement bound m, coefficient exponent C_1, and depth $J \ge m$. Write

$$M = J + 1, \qquad t_* = m + 1, \qquad n_* = 2\lceil L^{.52} \rceil, \qquad B_0 = \lceil L^2 \rceil, \qquad q = \frac{1}{256}, \qquad \varrho = \frac{1}{2}.$$

Every pattern has one unshared slot in a low group and at most t_* unshared slots in a high group. Its ordered free labels retain their original laws $\mu_g(p) = 1/(V_g p)$, including repeated values, and $\sum_\nu \|c_\nu\|_\infty \le L^{C_1}$.

The permutation group is the full direct product

$$\Pi = \prod_{g \in \mathcal{G}} \mathfrak{S}_M.$$

Source and target permutations are two independent uniform elements of Π. Equivalently, if $S_\Pi = |\Pi|^{-1} \sum_{\pi \in \Pi} U_\pi$ is the orthogonal averaging projection on label arrays, the physical operator is $S_\Pi A_{\mathrm{raw}} S_\Pi$. The ideal operator uses the corresponding product over $\mathcal{B}$ only. There is no permutation shared by different groups. In particular, selecting the omitted low slot produces M^{s_S} equally weighted choices.

For each *high* prime $p \in \bigcup_{g \in \mathcal{B}} \mathcal{P}_g$ put

$$d_j = \mathbf{1}_{j < n_*} + \mathbf{1}_{j > 0}, \qquad \eta_j = 1 - q^{d_j/2}, \qquad b_p = 1 - \frac{1}{p} \sum_{j=0}^{n_*} \eta_j, \qquad m_p = \frac{1}{p b_p}, \qquad \nu_g(p) = \frac{m_p}{V_g}.$$

Only the high block is tilted. In particular,

$$\Lambda = \max_{p \in \bigcup_{g \in \mathcal{B}} \mathcal{P}_g} b_p^{-1}, \qquad \beta = \prod_{g \in \mathcal{B}} \prod_{p \in \mathcal{P}_g} b_p.$$

For sufficiently large L,

$$1 \le \Lambda \le 1 + \frac{2(n_* + 1)}{P_*}, \qquad 1 \le N_g := \sum_{p \in \mathcal{P}_g} \nu_g(p) \le \Lambda, \qquad \nu_g(p) \le \frac{4}{P_*}. \tag{6.4}$$

These inequalities are not asserted for low primes.

The memory alphabet is

$$\mathcal{Z} = \bigsqcup_{g \in \mathcal{B}} \{(p, r) : p \in \mathcal{P}_g,\ r \in \mathbb{Z}/p\mathbb{Z}\}.$$

A memory s is a multiset on $\mathcal{Z}$, with total size $|s|$. Let $Z_g(s) = \sum_{p \in \mathcal{P}_g} s(p, 0)$ and $Z(s) = \sum_g Z_g(s)$. Its factorial mass is

$$\pi(s) = \prod_{g \in \mathcal{B}} \prod_{p \in \mathcal{P}_g} \prod_{r \bmod p} \frac{\nu_g(p)^{s(p,r)}}{s(p, r)!}.$$

The state also includes the small physical vertex $x_S = (z, \mathbf{p}_S)$ and the full ordered high active array $\mathbf{a}_B$. With $\tau_S = \prod_{g \in \mathcal{S}} V_g^{-M}$ its mass is

$$m_X(x_S, \mathbf{a}_B, s) = \tau_S \left(\prod_{g \in \mathcal{B}} \prod_{i=1}^{M} \nu_g(a_{g,i}) \right) \pi(s). \tag{6.5}$$

Endpoint active arrays are distinct within each group. Intermediate auxiliary active labels used to factor an edge range over the full product alphabet. The memory cutoff is imposed after every store, translation, and retrieval.

All row kernels act by ordinary sums. Thus their weighted adjoint is

$$K^*(x, y) = \frac{m_X(y)}{m_X(x)} \overline{K(y, x)}. \tag{6.6}$$

The elementary identity $\pi(s + [u])(s(u) + 1) = \nu_{g(u)}(p(u))\pi(s)$ is the source of every retrieval multiplicity.

6.3.2 Store, retrieval, and phases on a complete edge

We give the local calculation that is needed for phase-modified edges; compare [7, (5.20)–(5.34)]. For one high group, suppress its subscript and put $e(a) = (a, 0)$. One slot has the store kernel

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

The other active slots are fixed. Formula (6.6) and the factorial identity give

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

A retrieval therefore either samples a fresh target label with mass $\nu(a')$, or removes a pending copy with its original multiplicity.

For the $t \le t_*$ unshared slots of a group, in increasing original order $i_1, \dots, i_t$, the source block is

$$\mathcal{S} = D_\varrho S_{i_1,1,b_1} \cdots S_{i_t,1,b_t}, \qquad (D_\rho f)(\mathbf{a}, s) = \rho^{Z(s)} f(\mathbf{a}, s).$$

The target block is its weighted adjoint and processes the slots in decreasing original order. There is one source damping factor and one final target damping factor per group, not one factor per slot. For $0 < \alpha < 1$ the exact identities

$$S_{i,\alpha,b\alpha^{-r}} D_{\alpha^r} = D_{\alpha^{r+1}} S_{i,1,b}, \qquad \prod_{j=1}^{t} S_{i_j,\alpha,b_j\alpha^{-(t-j)}} = D_{\alpha^t} \prod_{j=1}^{t} S_{i_j,1,b_j}$$

follow by considering separately the drop and store alternatives. Set $\alpha = 1 - 1/(2t_*)$; then $\alpha^{t_*} \geq 1/2 = \varrho$. Consequently the once-damped block is the displayed distributed product preceded by a diagonal contraction.

Use the Schur weight $\mathfrak{w}(s) = 16^{|s|}$. A slot with store coefficient one has weighted row constant at most $N_g(1+16)$ and weighted reversed-column constant at most $N_g + [16(1-\alpha)]^{-1}$. The second estimate uses $h\alpha^{h-1} \leq (1-\alpha)^{-1}$. The distributed store coefficients are at most $2\alpha^{-t_*}$. Thus the whole source block has row and reversed-column constants R_0^t and C_0^t, where

$$R_0 = 2\alpha^{-t_*}\Lambda(1+16), \qquad C_0 = 2\alpha^{-t_*}\left(\Lambda + \frac{1}{16(1-\alpha)}\right).$$

Its adjoint has these constants interchanged.

Translation sends (p, r) to $(p, r + \Delta)$ and preserves π, $|s|$, and the Schur weight. The normalized memory edge is $N_g^{-t}\mathcal{S}_s U_\Delta \mathcal{S}_r^*$. The factor N_g^{-t} has a precise purpose: the auxiliary active label introduced at each stored slot is summed over the full law ν_g, contributing N_g. It cancels before any estimate and remains correct with arbitrary coefficients on the final endpoint.

For an original pattern and free tuple, the displacement is $\Delta = kD_\nu$. It is fixed while processing the unshared slots, because D_ν uses only shared source labels and free labels. The store and retrieval coefficients in a forward edge are

$$b_s(p) = \mathbf{1}_{p\nmid D_{\nu,B}}, \qquad b_r(p) = V_g^{-1}\mathbf{1}_{p\nmid D_{\nu,B}}.$$

Both are at most two. A newly stored allowed prime is translated to a nonzero residue, since $p \nmid k$ and $p \nmid D_{\nu,B}$; it cannot be retrieved at that edge. The scalar multiplying the product of memory kernels uses the low-block damping

$$d_{q,S}(z, \mathbf{p}_S) = q^{(\sum_{g\in\mathcal{S}} \omega_g(z) - Ms_S)/2}$$

and is exactly

$$\frac{d_{q,S}(x_S)d_{q,S}(y_S)}{|\Pi|^2}\left(\prod_{g\in\mathcal{S}} V_g^{-1}\right)\left(\prod_g \prod_{i\in C_{\nu,g}} \mu_g(r_{g,i})\right) c_\nu(\mathbf{a}_B, \mathbf{a}'_B, \mathbf{r}_B) D_\nu^{i\zeta} e(\theta k D_\nu)\mathcal{I}_\nu. \tag{6.7}$$

Here $\mathcal{I}_\nu$ retains the integer edge equation, shared labels, endpoint distinctness, and exclusions of free labels. The inverse high-group normalization has already entered the fresh mass ν_g or the retrieval coefficient V_g^{-1}; it must not be inserted again in (6.7). Reversing an edge conjugates this scalar, swaps its endpoints and the two coefficient families, and replaces Δ by $-\Delta$. This is precisely the mass-weighted adjoint, not an unweighted transpose on memory states.

Positive restriction to the common cutoff preserves the Schur bounds. For one high group the two edge constants are at most $(R_0C_0)^t$. Since $\Lambda \leq 2$ and $\alpha^{-t_*} \leq 2$ eventually,

$$R_0C_0 \leq 68(t_* + 16).$$

The low target completion has row and reversed-column constant at most

$$C_S^{\mathrm{Schur}} = \prod_{g\in\mathcal{S}} \frac{1}{(1-\sqrt{q})V_g} \le L^{1/5}.$$

Indeed, if h primes remain eligible beyond the fixed shared labels, their contribution is at most $h(\sqrt{q})^{h-1}/V_g$. The original free laws have mass one and the two permutation averages have mass one. Taking the coefficient budget once, these calculations give, even when every original edge choice is multiplied by an arbitrary complex phase of modulus one,

$$\|E_{\mathrm{phase}}\| \le L^{C_\ell}, \qquad C_\ell = \max(0, C_1) + 1 + \frac{1}{5} + .07t_* \log\bigl(68(t_*+16)\bigr). \tag{6.8}$$

The exponent is fixed before J. Arbitrary phases need not factor through individual stores: their absolute kernel is dominated by the positive sum just bounded, in both the row and reversed-column tests.

6.3.3 The small norm of an unchanged complete edge

The finite completion and Fourier identities [7, (5.37)–(5.52)] apply to the present alphabets. For clarity, the two facts needed for that assertion are as follows. Completion of the one unshared low label has squared norm at most

$$C_{\mathrm{cond}} = \prod_{g\in\mathcal{S}} \frac{1}{(1-q)V_g} \le L^{1/5}.$$

The residue translations preserve the factorial mass (6.5), so the finite orbit Fourier decomposition is an isometry and its complete frequency blocks are the current ideal row operators. Changing μ_g to $\nu_g = \mu_g b_p^{-1}$ in the at most $(m+1)s_B$ resampled high coordinates costs at most $\Lambda^{(m+1)s_B} - 1$. Excluding collisions of a free high label with an endpoint label costs at most $2m(J+m+2)s_B/P_*$ by the union bound.

The additional store or retrieval part is controlled by the low omissions. Freeze the pattern, the free labels, the high source permutation, and the entire target permutation. Put

$$S_{\mathrm{low}} = \prod_{g\in\mathcal{S}} \prod_{i=1}^{M} p_{g,i},$$

and let U_{om} be the product of the omitted low source labels. Since $C_{\nu,g} = \varnothing$ for $g \in \mathcal{S}$, the displacement is $\pm k D_{\nu,B} S_{\mathrm{low}}/U_{\mathrm{om}}$, with $D_{\nu,B}$ fixed. A pending copy (p,h) can be retrieved with nonzero weight only if $p \nmid D_{\nu,B}$ and

$$hU_{\mathrm{om}} + (\pm k) D_{\nu,B} S_{\mathrm{low}} \equiv 0 \pmod{p}.$$

Here $p \nmid k$ by hypothesis and $p \nmid S_{\mathrm{low}}$ by the separation of the low and high prime sets. The constant term is nonzero modulo p. If $h = 0$ in $\mathbb{Z}/p\mathbb{Z}$, there is no solution; otherwise the congruence determines one residue class for U_{om}. The M^{s_S} omission products are distinct by unique factorization and satisfy

$$1 \le U_{\mathrm{om}} \le e^{s_S L^{.36}} < P_* < p \qquad (p \text{ a high prime}).$$

Each pending copy therefore permits at most one equally weighted omission product. With at most B_0 pending copies, averaging the full product $\prod_{g\in\mathcal{S}} \mathfrak{S}_M$ gives proportion at most B_0/M^{s_S}. Newly stored copies have $h = 0$ and cannot be retrieved on the same edge. Weighted

Schur, with the positive memory constants above, gives the promotion factor $\sqrt{B_0/M^{s_S}}$. The mass-weighted adjoint treats the reverse correction in the same way.

Writing $\epsilon = \sup_\Theta \|\mathcal{T}_\zeta(\Theta)\|$ and $K = \sum_\nu \|c_\nu\|_\infty$, the resulting complete signed edge is bounded by

$$\|E\| \le C_{\mathrm{cond}} \left\{ \epsilon + K(\Lambda^{(m+1)s_B} - 1) + K\Lambda^{(m+1)s_B} \frac{2m(J+m+2)s_B}{P_*} \right\} + 4KC_S^{\mathrm{Schur}}[68(t_*+16)]^{t_* s_B} \sqrt{\frac{B_0}{M^{s_S}}}. \tag{6.9}$$

These are the finite completion identities and the bounds in [7, Lemma 5.3], with the current laws in every sum. The checks above account for their arithmetic uses of the alphabets; they do not use their contiguity.

If $\epsilon \le L^{-(G+3)}$, a sufficiently large fixed J makes $\|E\| \le L^{-G}$. For example it is enough that

$$.0225 \log(J+1) \ge G + C_1 + \frac{1}{5} + .07 t_* \log(68(t_*+16)) + 4.$$

After J is fixed, the collision and tilt terms are smaller than every fixed logarithmic power. The same bound holds for the reverse edge. This small norm concerns the complete signed sum of patterns. It is not asserted for its positive majorant.

6.3.4 The rank expansion and its boundary normalization

The elementary residue identities behind the expansion are

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

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

where $H_j(r) = \mathbf{1}_{r + h_j = 0 \bmod p}$, $A \neq \varnothing$, and $a_0 \in A$. They are identities on a single residue ring; see [7, (5.15)–(5.16)]. The first follows by expanding $\prod(1 - x_j)$ into its constant, one-point, and consecutive surviving two-point terms; the second fixes $r = -h_{a_0}$. Processing these identities in chronological order gives stores, translations, retrievals, and ghost visits. The initial and final pending memories are empty.

Potential birth addresses form the disjoint union

$$\mathcal{I} = (\mathcal{B} \times [M]) \sqcup ([0, n_* - 1] \times \mathcal{B} \times [M]) \sqcup ([0, n_*] \times [0, B_0 - 1]).$$

They describe initial active labels, fresh active labels on edges, and the ordered fresh batches at ghost visits. An unused address a receives its own symbol $\perp_a$. For a history h, write $\widehat{p}_h(a)$ for its tagged prime or unused symbol, and set

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

The determinant of the equality matrix gives

$$\sum_{\sigma \in \mathfrak{S}(\mathcal{I})} \mathrm{sgn}(\sigma) \Delta_\sigma(h) = \mathbf{1}_{\text{all performed birth primes are distinct}}. \tag{6.10}$$

Unused symbols supply identity rows. This identity imposes global distinctness after, and not before, the signed sum.

Put $Q_S = \prod_{g\in\mathcal{S}} \prod_{p\in\mathcal{P}_g} p$. For $0 \le z < Q_S$, the initial coefficient and the terminal coefficient in the row expansion are

$$L_z(x) = \mathbf{1}_{\mathrm{pos}(x)=z}\mathbf{1}_{s=\varnothing} \prod_{g\in\mathcal{B}} \prod_{i=1}^{M} \nu_g(a_{g,i}), \qquad R_z(x) = \mathbf{1}_{\mathrm{pos}(x)=z}\mathbf{1}_{s=\varnothing}.$$

Let $W_\sigma(z)$ denote the sum of original history weights times Δ_σ, with these coefficients. The normalized quantity is

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

The terminal active array is freely summed; it is not identified with the initial array.

Let C_z count the ordered distinct low arrays dividing z, and let

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

In the Hilbert convention (6.5), the initial vector is $\overline{L_z/m_X} = \tau_S^{-1} R_z$, rather than L_z. Its squared norm is $\tau_S^{-1} S_0 C_z$; the squared norm of the terminal vector is $\tau_S S_0 C_z$. Their norm product is $S_0 C_z$. The Chinese remainder theorem gives the exact boundary identity

$$\frac{\tau_S}{Q_S} \sum_{z \bmod Q_S} C_z = \prod_{g\in\mathcal{S}} \sum_{\substack{p_1,\dots,p_M\in\mathcal{P}_g \\ \text{pairwise distinct}}} \prod_{i=1}^{M} \mu_g(p_i) \le 1. \tag{6.11}$$

Moreover

$$S_0 \le \left(1 + \frac{2(n_*+1)}{P_*}\right)^{M s_B} \le 2$$

for each fixed J and all sufficiently large L. Since $0 < \beta \le 1$, the normalized product of boundary norms is at most two. This is the normalization of [7, (5.55)–(5.57)].

6.3.5 High rank: prescribed fresh draws

Order all birth addresses chronologically, placing initial addresses first. In each nontrivial cycle of σ, choose the first address as pivot and prescribe the value at every later address to equal its pivot. There are exactly $r(\sigma)$ prescribed addresses. Every moved address must be performed, since $\perp_a$ occurs only at a. A cycle with two initial addresses contributes zero, by distinctness of the initial active labels and disjointness of the groups. Thus each nonzero prescription is a fresh draw on an edge or at a ghost, and its pivot has already been exposed.

We record enough of the positive row calculation to make its uniformity in the current alphabets explicit. Put $v = 16$, $\alpha_h = 1 - 1/[2(t_*+1)]$, and

$$H_{t_*} = 2 + \frac{2\alpha_h^{-t_*}}{16(1-\alpha_h)}, \qquad K_{t_*} = [33\max(1, H_{t_*})]^{t_*}, \qquad \delta_* = \frac{64}{P_*}.$$

Process target slots in decreasing original order. Applied to $\alpha_h^{(k+1)Z_g(s)} v^{|s|}$, a free target draw costs at most two, while a retrieval costs at most $2\alpha_h^{-k}/[16(1-\alpha_h)]$. These bounds follow from (6.4) and $h\alpha_h^{h-1} \le (1-\alpha_h)^{-1}$. A prescribed slot can only use its fresh alternative and costs at most $4/P_* \le \delta_*$. A prescription at a slot which is absent or retrieved contributes zero. Summing source-store masks costs at most $(1+2v)^{t_g}$ in group g.

The low-label completion is at most $(47/15)^{s_S}$, and free-label and permutation sums retain total mass one. Consequently a complete absolute edge row, including all its original guards and any record of earlier exposed primes, has the bound

$$\sum_{\text{edge choices}} |\text{coefficient}|\,\mathbf{1}_{\text{prescriptions}}\,16^{|s'|} \le L^{C_e}\delta_*^{r_i^E}\,16^{|s|},$$
$$C_e = \max(0, C_1) + 1 + .09\log(47/15) + .07\log K_{t_*}. \tag{6.12}$$

The proof is the finite conditional sum in [7, (5.63)–(5.71)]. Its displacement and shared labels are fixed before the target fresh draws, in either orientation. No independence of the other target choices is assumed.

At ghost visit j, a retained multiplicity $h_u \le s_u$ contributes $\binom{s_u}{h_u}A_j(u)^{h_u}D_j(u)^{s_u-h_u}$, and an ordered fresh batch $(u_1,\dots,u_b)$ contributes $b!^{-1}\prod_{\ell=1}^{b} C_j(u_\ell)\nu_{g(u_\ell)}(p(u_\ell))$, where

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

The factorial of the ordered batch is retained before any prescription is imposed. A prescribed fresh prime, including its factor v, has total residue mass at most $32m_p \le 64/P_*$. An unrestricted fresh coordinate has total weighted mass at most $32\sum_p m_p$. Retentions satisfy

$$\prod_u (A_j(u)v + |D_j(u)|)^{s_u} \le v^{|s|}.$$

Indeed the factor is v off residue zero and is at most $4(1-\eta_j)v + 2\eta_j \le 6$ at zero. Summing batch lengths with their factorials therefore gives

$$\sum_{\text{ghost choices}} |\text{coefficient}|\,\mathbf{1}_{\text{prescriptions}}\,16^{|s'|} \le e^{32\sum_p m_p}\delta_*^{r_j^G}\,16^{|s|} \le L^7\delta_*^{r_j^G}\,16^{|s|}. \tag{6.13}$$

When a pivot occurs earlier in the same ordered batch, integrate the coordinates backwards, conditioning on their full preceding record. This proves the same bound without a sum over possible records. These are the original factorial calculations in [7, (5.72)–(5.78)].

Iterating (6.12) and (6.13), and using the normalized initial mass in (6.11), proves

$$|w_\sigma| \le 2L^{C_h(n_*+1)}\left(\frac{64}{P_*}\right)^{r(\sigma)}, \qquad C_h = C_a + 8, \quad C_a \ge C_e. \tag{6.14}$$

The exponent is independent of J; only the final threshold depends on the fixed depth.

6.3.6 Low rank: local phases and the signed operator product

For every address a and every tagged prime or unused symbol u, choose independent uniform signs $\epsilon_a(u) \in \{-1, 1\}$. For fixed history values,

$$\begin{aligned}\Delta_\sigma(h) &= \mathbb{E}_\epsilon \prod_{a\in\mathcal{I}} \epsilon_a(\widehat{p}_h(a))\epsilon_a(\widehat{p}_h(\sigma a)) \\ &= \mathbb{E}_\epsilon \prod_{a\in\mathcal{I}} \epsilon_a(\widehat{p}_h(a))\epsilon_{\sigma^{-1}a}(\widehat{p}_h(a)).\end{aligned} \tag{6.15}$$

The second product is local in chronological time: at a performed address its factor is evaluated on the prime drawn there; at an unused address it is a fixed sign evaluated on $\perp_a$. At a fixed point of σ the factor is identically one.

A cycle of length $d \geq 2$ has $d \leq 2(d-1)$ addresses. Hence the support of σ has at most $2r(\sigma)$ addresses, and at most $2r(\sigma)$ edge operations acquire a phase. Initial-address phases only multiply the initial vector by a function of modulus one. Ghost phases are handled by the following two-sided estimate.

For a single memory coordinate u with $\lambda_u = \nu_{g(u)}(p(u))$, the absolute ghost kernel is

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

This formula follows by collecting the ordered fresh batches with their original factorial. The factorial masses give

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

The weighted row sum with $16^{|s|}$ is bounded by $\exp(16 \sum_u |C_j(u)|\lambda_u)$, because $A_j + |D_j|/16 \leq 1$. The reversed-column sum interchanges C_j, D_j and has the analogous bound, since $A_j + |C_j|/16 \leq 1$. Both inequalities hold at zero residue ($A_j \leq 1/4$, $|D_j| \leq 2$, $|C_j| \leq 3$), and trivially elsewhere. The other state masses cancel because the ghost leaves its active array and small vertex fixed. Schur's inequality therefore gives, with arbitrary local phases and with the common cutoff,

$$\|G_{j,\mathrm{phase}}\| \leq \exp\left(48 \sum_{g \in \mathcal{B}} \sum_{p \in \mathcal{P}_g} \nu_g(p)\right) \leq e^{96 s_B} \leq L^8. \tag{6.16}$$

Deleting positive terms for a cutoff preserves these bounds.

Choose $C_a \geq \max(C_e, C_\ell, 0)$ before J. For each fixed assignment of signs in (6.15), an unaffected edge is the original complete signed operator E or E^* and has norm at most L^{-G}. Affected edges have norm at most L^{C_a} by (6.8). It would be incorrect to replace every edge by its positive majorant: that would remove the factor L^{-Gn_*}. Multiplying the operator norms in their chronological order, and then using the boundary normalization, gives

$$|w_\sigma| \leq 2L^{8(n_*+1)} L^{-Gn_*} \left(L^{G+C_a}\right)^{2r(\sigma)} \leq L^{(-G+8)n_* + 9 + 2r(\sigma)(G+C_a)}. \tag{6.17}$$

Averaging the signs does not alter the bound. This proves the low-rank estimate, with its original signed operators, for the current prime alphabets.

6.3.7 Summing the ranks

For each fixed J, $|\mathcal{I}| \leq L^4$ eventually. There are at most $|\mathcal{I}|^{2r} \leq L^{8r}$ permutations of rank r, since a rank-r permutation is a product of r transpositions. Fix $E_0 > 0$, take $G = E_0 + 9$, and put

$$r_0 = \left\lceil \frac{2(C_h + E_0 + 5)(n_* + 1) \log L}{L^{.39}} \right\rceil.$$

For $r \geq r_0$, use (6.14) and $L^8(64/P_*) \leq e^{-L^{.39}/2}$. The geometric sum is at most $L^{-E_0 n_*}$. For $r < r_0$, use (6.17); its sum is at most

$$r_0 L^{(-G+8)n_* + 9 + \{8 + 2(G+C_a)\} r_0} \leq L^{-E_0 n_*},$$

because $r_0/n_* \to 0$ and $G - 8 = E_0 + 1$. Consequently

$$|\mathfrak{p}_{n_*}| \leq 2L^{-E_0 n_*}. \tag{6.18}$$

6.3.8 The uncapped memory and its discarded absolute mass

Let $\mathfrak{m}_{n_*}$ be the independent-residue moment before the memory cutoff. The elementary identities above identify it with the uncapped lifespan expansion, with the same baseline β, empty boundary memories, and original guards. There is a direct bound on the absolute mass removed by the cutoff. The following birth-count calculation makes the post-summation bound in [7, (5.84)] explicit. For this bound keep the indicator of global birth distinctness in the lifespan expansion, take absolute values there, and only then majorize that indicator by one. Do not first expand it into the determinant sum and sum $|w_\sigma|$.

Call a *ghost birth* an entry in an ordered fresh ghost batch, counted with multiplicity. This is a cumulative count over the whole history, not the number of particles currently pending. Initial active labels are not pending particles. Each of the n_* edges stores at most $t_* s_B$ particles; retrievals and retentions produce no new particles. Thus a history whose pending memory ever exceeds B_0 has at least

$$b_{\mathrm{ghost}} \geq B_0 + 1 - n_* t_* s_B$$

ghost births. This includes an overflow immediately after a store, even when a later retrieval would reduce the memory.

Multiply the absolute coefficient of each ghost birth by two. In (6.13), this changes only the fresh intensity, from $32 \sum_p m_p$ to $64 \sum_p m_p \leq 192 s_B$. The retention inequality and the edge estimate (6.12) remain unchanged. Both estimates are uniform in the common cutoff. One may first use an arbitrary larger finite cutoff and then pass to the nonnegative sum by monotone convergence; the bounds supply its finite total mass. The original guards can be retained, or removed only when enlarging this positive majorant. The normalized initial mass is at most two. Markov's inequality for $2^{b_{\mathrm{ghost}}}$ therefore proves

$$|\mathfrak{m}_{n_*} - \mathfrak{p}_{n_*}| \leq 2\, 2^{-B_0-1+n_* t_* s_B} L^{C_e n_*} \exp\bigl(192(n_*+1)s_B\bigr) \leq L^{-E_0 n_*} \tag{6.19}$$

for every fixed E_0, once L is sufficiently large. Here $B_0 \asymp L^2$ whereas $n_* s_B = O(L^{.52} \log L)$. In particular this implies the coarser source bound $2^{-B_0-1+n_* M s} \exp(O_{m,C_1,J}(n_* L \log L))$. This estimate is taken after the absolute weights of all histories have been summed. It does not infer a tail estimate merely from the signed cancellation in (6.18).

6.3.9 Physical counting versus independent residues

We now justify the counterpart of [7, (5.83)], keeping track of the subsequent path sum. Put $U = \lfloor e^L \rfloor$, $H = n_* U^{2Ms+1}$, and $\mathcal{W} = [-H, -H + 2Q_S) \cap \mathbb{Z}$. The group bounds imply

$$\log Q_S \geq \tfrac{1}{2} e^{L^{.27}} L^{.27}, \qquad 2H \leq Q_S, \qquad 2X \leq Q_S \quad (e^{.75L} \leq X \leq e^{1.25L}).$$

Every path of n_* edges from $[X, 2X)$ stays in $\mathcal{W}$. Its displacement at any edge is at most U^{2Ms+1}.

Fixed labels and the relative Bonferroni error. Fix every pattern, label, free tuple, and endpoint permutation along a closed path. Incompatible active congruences contribute zero. Otherwise the active congruences fix one residue modulo a squarefree q_{act} with

$$\log q_{\mathrm{act}} = O_J(n_* s L^{.46}) = O_J(L^{.98} \log L) = o(L).$$

Keep their deterministic damping as a factor $D_{\mathrm{act}} \geq 0$. For a prime not in the active labels, combine all its visits into $1 - g_p(z)$, with $0 \leq g_p \leq 1$. It is supported on at most $n_* + 1$ residue

classes modulo p. Writing $a_p = p^{-1}\sum_{r \bmod p} g_p(r)$ gives

$$a_p \le \frac{n_*+1}{p} = o(1), \qquad \sum_p a_p = O(n_* \log L).$$

The independent-residue mean of the damping factor is $D_{\mathrm{act}} q_{\mathrm{act}}^{-1} \prod_p (1-a_p)$, and

$$\prod_p (1-a_p) \ge \exp(-C n_* \log L).$$

Use two adjacent Bonferroni truncations at order $K_0 = \lceil C_0 n_* \log L \rceil$. Their expected gap is at most $\left(\sum_p a_p\right)^{K_0}/K_0!$. Given any fixed $D > 0$, choosing C_0 sufficiently large makes this gap, relative to $\prod_p (1-a_p)$, at most L^{-Dn_*}. The constants are uniform over the fixed path labels.

The absolute mass over all label paths. To absorb that relative error, define a positive periodic operator $\mathcal{A}_+$ on positions modulo $\mathcal{Q} = \prod_{p \in \mathcal{P}} p$. It uses the physical kernels with each original pattern coefficient replaced by its absolute value and each phase by one, before the patterns, free labels, and independent permutations are summed. The coefficient budget is unchanged. As $\mathcal{Q} \ge Q_S > 2H$, a path of n_* edges which closes modulo $\mathcal{Q}$ has zero integer displacement. Its averaged root moment is therefore exactly the absolute independent-residue path sum $\mathfrak{m}_{n_*}^{\mathrm{abs}}$.

Here the Schur bound is particularly simple. A physical source has at least M distinct primes in each group, so its damping is at most one. For a fixed pattern and free tuple, the target position is fixed by the shared labels. If it has $M + j$ distinct primes in group g, at most $j + t_g$ labels remain beyond its $M - t_g$ fixed shared labels. The sum over its ordered unshared target labels, including its damping and normalization, is at most

$$\frac{(j+t_g)_{t_g} q^{j/2}}{V_g^{t_g}} \le \frac{2^{t_g} t_g!}{(1-\sqrt{q})^{t_g+1}}.$$

The estimate follows by bounding the supremum by the series $\sum_{j \ge 0} (j+t)_t (\sqrt{q})^j = \frac{t!}{(1-\sqrt{q})^{t+1}}$. All other guards only remove terms. Put

$$C_* = \max_{1 \le t \le t_*} \frac{2^t t!}{(1-\sqrt{q})^{t+1}}, \qquad C_{\mathrm{abs}} = C_1 + .16 \log C_* + 1.$$

The free laws and permutation averages have mass one, so every row sum is at most $L^{C_1} C_*^s \le L^{C_{\mathrm{abs}}}$. The same calculation for the reversed edge gives its column sum. The physical mass $\tau = \prod_g V_g^{-M}$ is constant on vertices, so Schur's inequality yields $\|\mathcal{A}_+\| \le L^{C_{\mathrm{abs}}}$.

Let v_z be the full position-fibre indicator in this periodic graph and let C_z^{all} count its ordered distinct label arrays. The Chinese remainder theorem gives

$$\frac{1}{\mathcal{Q}} \sum_{z \bmod \mathcal{Q}} \|v_z\|^2 = \frac{\tau}{\mathcal{Q}} \sum_{z \bmod \mathcal{Q}} C_z^{\mathrm{all}} = \prod_g \sum_{\substack{p_1,\dots,p_M \in \mathcal{P}_g \\ \text{pairwise distinct}}} \prod_{i=1}^{M} \mu_g(p_i) \le 1.$$

Consequently

$$\mathfrak{m}_{n_*}^{\mathrm{abs}} = \frac{1}{\mathcal{Q}} \sum_z \langle v_z, (\mathcal{A}_+ \mathcal{A}_+^*)^{n_*/2} v_z \rangle \le L^{C_{\mathrm{abs}} n_*}. \tag{6.20}$$

This calculation retains the two independent root label arrays. The exponent C_{abs} is fixed before J. Choosing $D > E_0 + C_{\mathrm{abs}} + 2$ in the Bonferroni estimate therefore absorbs the sum of its relative errors over all paths. There is no signed-rank or physical-to-residue estimate in this positive periodic bound.

Rounding errors in the retained CRT terms. Each g_p can be written as a sum of at most $n_* + 1$ residue indicators, with coefficients in $[0, 1]$, by combining visits to the same residue class. It must be expanded in this form; expanding separately into all subsets of visits is unnecessary. Every retained term is consequently one CRT condition with modulus of logarithm

$$O_{J,E_0}(n_* \log L\, L^{.46}) = o(L).$$

Counting that condition in $[X, 2X)$ has error at most one. There are at most $O_J(n_* s)$ label slots in a path, every label is at most $e^{L^{.46}}$, and the additional Bonferroni terms choose at most K_0 primes and K_0 residue classes. The logarithm of the number of these choices is $O_{J,E_0}(n_* \log L\, L^{.46})$. Arbitrary pattern multiplicity is handled by its coefficient sum: at n_* edges it costs at most $L^{C_1 n_*}$. The remaining normalization factors, inverse damping factors factored from active labels, and permutation sums cost $L^{O_J(n_*)}$ or less. It follows that the sum of all absolute rounding errors is

$$\exp(O_{J,E_0,m,C_1}(n_* \log L\, L^{.46})) = \exp(o(L)).$$

After division by $X \geq e^{.75L}$ this is smaller than $L^{-E_0 n_*}$. Invertibility of k at every group prime makes the residue changes in the path identities bijective.

Combining the relative and rounding estimates, with their precisions increased by fixed margins, proves

$$\left| \sum_{z \in [X,2X)} \langle u_z, (\mathcal{A}\mathcal{A}^*)^{n_*/2} u_z \rangle - X \mathfrak{m}_{n_*} \right| \leq X L^{-E_0 n_*}. \tag{6.21}$$

The constants and thresholds are uniform over all current prime alphabets and coefficient families satisfying their fixed data.

6.3.10 Conclusion and order of choices

For the desired moment exponent A, choose $E_0 = A + 1$ and $G = E_0 + 9 = A + 10$. Fix m, C_1, then the constants C_ℓ, C_e, C_a, C_h, and choose J large enough for (6.9) at ideal precision $L^{-(G+3)} = L^{-(A+13)}$. The positive exponent C_{abs} and the auxiliary Bonferroni precision can also be fixed before J. After this fixed depth, choose the threshold in L. These choices are uniform before the prime family, coefficients, phases, translation, and source scale. Equations (6.18), (6.19), and (6.21) yield

$$\sum_{z \in [X,2X)} \langle u_z, (\mathcal{A}\mathcal{A}^*)^R u_z \rangle \leq 4X L^{-(A+1)2R} \leq X L^{-2AR}.$$

The row expansion closes the position and sums its two endpoint label arrays independently, with the initial vertex mass once; it is exactly the physical root moment in (6.3).

The proof uses the ideal-norm premise and the explicit finite memory identities of the cited construction. It uses neither an orthogonality assertion about the eventual arithmetic sources nor an absolute-value substitute for the low-rank signed product. The additional analytic inputs needed for the marked correlation are treated separately in the arithmetic comparison.

6.4 The prime estimate used for a coloured group

The arithmetic input for the groups comes from the zero-free half-plane in [8, Theorem 1.1], in the form imported and used in [7, Theorem 3.2, Lemma 3.5 and Corollary 3.6]. It states that, for every Dirichlet character and every modulus, the continued L-function has no zero

in $\Re s > 7/8$, with the principal pole at $s = 1$ excluded. The resulting prime estimate in the range needed here is

$$\left| \sum_{a<p\le b} \psi(p) p^{it} \right| \ll a \exp(-L^{.27}/512), \quad e^{L^{.27}} \le a < b \le 2a, \tag{6.22}$$

uniformly for nonprincipal characters modulo $q \le e^{2L}$ and $|t| \le e^{3L}$. The upper endpoints used here are at most $2\exp(L^{.46})$. One obtains this form from the reciprocal prime estimate $O((1+\log((q+2)(|t|+3)))a^{-1/256})$ of [7, Lemma 3.5]; partial summation gives the unweighted form recorded in its Corollary 3.6. We use this as a specified external analytic input; the proof of the zero-free theorem is not reproduced.

For a regular colouring, split its character sum into s_i^{-1} times the sum over its prime block and the centred colour discrepancy (2.10). Omitting the primes dividing T costs at most $O(L)$ in an unweighted sum. Intersecting the dyad with the block still gives an interval, so (6.22) applies. The discrepancy is at most $a\exp(-L^{.27}/100)$ up to an absolute factor. Their sum therefore has a stretched exponential saving, and in particular is $O_A(aL^{-A})$ for every fixed $A > 0$, uniformly over all regular colourings. Subdivision and partial summation give the corresponding reciprocal estimates. This verifies the prime-sum premise at its point of use; concentration by itself would only compare a group to its parent block.

6.5 Producing the ideal coefficients for dispersed groups

For the actual coefficient construction one also needs the character comparison on each group. A sufficient condition is that for every fixed $D, C > 0$,

$$\sup_{\substack{\psi \ne \psi_0 \pmod q,\ q \le L^C \\ I \text{ an interval}}} \left| \sum_{p \in I \cap \mathcal{P}_g} \frac{\psi(p)}{p} \right| = O_{D,C}(L^{-D}), \tag{6.23}$$

uniformly in g and the prime family under consideration. The good colourings constructed earlier satisfy this condition by their uniform prime-sum estimate and partial summation.

All product cells must be rebuilt with the current laws μ_g. For a selected coordinate set A, let $j(v) = \lceil \log(v)/h \rceil$ and $\nu_{A,j} = \Pr(j(P_A) = j)$. Define

$$\mathcal{K}_A(D, X) = \frac{\mathbf{1}_{j(D)=j(X)} \mathbf{1}_{\nu_{A,j(X)} \ge \xi}}{\nu_{A,j(X)}} \sum_{\chi \in \mathcal{C}(Q_0)} \overline{\chi(D)} \chi(X),$$

where $Q_0 = \lceil L^{q_*} \rceil$, $h = L^{-h_*}$, $\xi = L^{-x_*}$, and $\mathcal{C}(Q_0)$ is the set of primitive characters of conductor at most Q_0. A zero cell contributes zero. The denominator is the actual cell mass, not the mass of an interval envelope or of a previous partition.

More explicitly, split $\mathcal{B} = \mathcal{B}_1 \sqcup \mathcal{B}_2$ into blocks whose sizes differ by at most one, and let I_j consist of the first m slots in the groups of $\mathcal{B}_j$. The raw pattern has $C_{0,g} = \varnothing$ and $c_0 = 1$. For $\varnothing \ne A \subseteq I_1$, $\varnothing \ne B \subseteq I_2$, the pattern replaces the slots of $A \cup B$ and has coefficient

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

In particular, the first kernel reads the target labels and the second reads the source labels. Each free tuple retains its original product law, including repetitions.

When all selected labels except one prime are fixed, a cell condition restricts that prime to the intersection of its group with an interval. Equation (6.23) bounds every off-diagonal

Gram entry. The diagonal is the actual cell mass, since all group primes exceed the relevant conductors. More explicitly, the unnormalized Gram matrix has entries

$$G_j(\chi,\chi') = \mathbb{E}[\mathbf{1}_{j(P_A)=j}\overline{\chi(P_A)}\chi'(P_A)], \qquad G_j(\chi,\chi) = \nu_{A,j}.$$

The quotient of two distinct primitive characters of conductor at most Q_0 is nonprincipal on their common modulus, which is at most Q_0^2. Thus (6.23) is invoked with $C \geq 2(q_*+1)$, after q_* has been chosen. If its off-diagonal entries have modulus at most $\rho_0 = L^{-D_0}$, the normalized Gram matrix on a retained cell has norm at most $1 + Q_0^2\rho_0/\xi$. Choosing D_0 sufficiently large makes the cell operator and its transpose have norm at most two, and gives

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

on retained cells. This also fixes the conjugation convention in the coefficient formula: the free product is conjugated in the kernel, and the endpoint product is not. For a selected tuple with product n, disjointness identifies the group of every prime, including its multiplicity. In each group there are at most $m!$ orders for these primes, and the normalizations cost at most 2^m. Since $s_B \leq .07\log L$,

$$\Pr(P_A = n) \leq L^{\kappa_m}/n, \qquad \kappa_m = .07\log(2^m m!).$$

For a test f supported on the product dyad $U \leq P_A < 2U$, Cauchy–Schwarz on each fibre gives

$$\sum_n \left|\mathbb{E}[f\,\mathbf{1}_{P_A=n}]\right|^2 \leq \max_{U\leq n<2U}\Pr(P_A=n)\,\mathbb{E}|f|^2 \leq (L^{\kappa_m}/U)\|f\|_2^2.$$

This is the minor-arc input in [7, Lemma 5.1, (5.7)–(5.8)]. On the major arcs the same proof discards cells of mass below ξ and reproduces characters on the other cells by the Gram estimate above. Indeed the logarithm of a product is $O_m(L^{.46}\log L)$, giving $O_m(1+h^{-1}L^{.46}\log L)$ cells. Within a retained cell two logarithms differ by at most h. These bounds give the same discarded mass and phase variation estimates as in that lemma. All four bounds use the actual coloured product law and are uniform in the partition. The subsequent projection argument is algebraic. Splitting $\mathcal{B}$ into two almost equal blocks, it bounds the residual ideal norm by

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

Choose m and A_0 after the desired precision, then q_*, h_*, x_*, and finally J. The coefficient count and cell bound give

$$\sum_\nu \|c_\nu\|_\infty \leq 1 + 2^{ms_B}Q_0^4\xi^{-2} \leq L^{C_1}, \qquad C_1 = .07m\log 2 + 4(q_*+1) + 2x_* + 1.$$

In particular C_1 is fixed before J. For example, eventually both block sizes are at least $.015\log L$, since $s_B \geq .035\log L$. For a desired ideal precision D, choose $.0075\log(m+1) > D+2$ and $A_0 > D + .07m\log 2 + 2$. Then the two projection terms and the summed cell errors are together at most L^{-D} for large L. With $D = A + 13$ the resulting pattern coefficients satisfy the ideal premise of G. This is the parametric counterpart of [7, Theorem 5.2]; composing it with G gives the constructed moment used here. The original Theorem 5.6 alone, whose groups are fixed intervals, would not justify this composition.

7 From the physical moment to the marked correlation

The physical moment controls a finite operator. To obtain a reflected correlation from it, one must identify the position functions to which the operator is applied, estimate all of its

comparison branches, and recover the original counting measure from its reciprocal source normalization. We give these steps explicitly. In particular, the unit restriction is part of the position function from the beginning; no density estimate for the nonunits of the total is used. We write $e(t) = \exp(2\pi i t)$.

Throughout this section the outer parameters are

$$L = \log N, \qquad Q = \exp(L^{.49}), \qquad L - L^{.20} \le \log T \le L, \qquad \vartheta = 1/256.$$

The prime families satisfy (6.1). The finite disjoint prime groups are indexed by $\mathcal{G} = \mathcal{S} \sqcup \mathcal{B}$. Low groups have their primes in $(\exp(L^{.27}), \exp(L^{.36})]$, and high groups in $(\exp(L^{.39}), \exp(L^{.46})]$. Every group prime is coprime to T, and

$$\tfrac{1}{2} \le V_g = \sum_{p \in \mathcal{P}_g} p^{-1} \le \tfrac{3}{2}, \qquad \mu_g(p) = \frac{1}{V_g p}, \qquad s = |\mathcal{G}| \le 2 \log L.$$

The group-prime character bounds are those already verified for regular colourings. In particular, for every nonprincipal character ξ modulo $q \le e^{2L}$, every $|\tau| \le e^{3L}$, and every interval $(a, b] \subseteq [a, 2a]$ meeting a group,

$$\left| \sum_{p \in \mathcal{P}_g \cap (a,b]} \xi(p) p^{i\tau} \right| \ll a \exp(-cL^{.27}) \tag{7.1}$$

with an absolute $c > 0$. All construction parameters below are fixed before the prime family, the total, and the readers are chosen.

Write n° for the part of n outside the union of the groups, $n^\flat = n_{\le Q}$, and

$$\omega_g(n) = \#\{p \in \mathcal{P}_g : p \mid n\}, \qquad \Omega_g(n) = \sum_{p \in \mathcal{P}_g} v_p(n).$$

Fix real readers a, b of absolute value at most one on, respectively,

$$Y_{\mathrm{src}}(n) = ((n^\circ)_{\le Q}, (\Omega_g(n))_g), \qquad Y_{\mathrm{tgt}}(n) = (n^\circ, (\Omega_g(n))_g).$$

We use a real source F, bounded by a fixed C_F, for which

$$F(dn) = \gamma(d) F(n) \quad (P^+(d) \le Q), \qquad \gamma \in \{1, \lambda\}. \tag{7.2}$$

The other source input is the one-variable estimate

$$\sum_{n \ge 1} F(n) c(n^\flat) w(n) \ll_{A,B,C} Z L^{-A} \tag{7.3}$$

for every fixed $A, B > 0$ and $C \ge 1$, uniformly when $e^{3L/4} \le Z \le e^{5L/4}$, $|c| \le L^B$, $\operatorname{supp} w \subseteq [Z/C, CZ]$, and $\|w\|_\infty + \int |w'| \le L^B$. The coefficient and the test may be complex. Proposition 5.3 supplies this estimate for each of F_λ, F_0, F_*, with the covariance in (7.2). Thus this is a one-variable input already proved for our three sources, rather than a reflected-correlation assumption.

Proposition 7.1 (Marked comparison with covariant sources, candidate graph argument)**.** *Use the constructed physical moment of Proposition 6.1 for the admitted prime families above. Let F be a real source satisfying (7.2) and (7.3), with $|F| \le C_F$. Put*

$$W(n) = \prod_{g \in \mathcal{G}} \frac{\omega_g(n) \vartheta^{\omega_g(n)-1}}{V_g},$$

where a factor is zero when its count is zero. For every fixed $A > 0$ and real $\chi \in C_c^\infty((0,1))$, the graph parameters and a threshold can be chosen so that

$$\frac{1}{T}\left|\sum_{\substack{0<n<T\\(n,T)=1}} \chi(n/T)W(n)W(T-n)F(n)a(Y_{\mathrm{src}}(n))b(Y_{\mathrm{tgt}}(T-n))\right| \le CL^{-A}\|a\|_\infty\|b\|_\infty. \quad (7.4)$$

The source reader has exactly the observation Y_{src}, and the target reader exactly Y_{tgt}; they are separate functions. The constant and threshold depend only on A, χ, C_F, the uniform constants in the one-variable and group-prime estimates, and the fixed moment construction. They are uniform in the readers, the admitted total T, and the admitted prime family, all of which may vary after the threshold. The same assertion is uniform over any family of sources for which the stated input bounds hold uniformly. The restriction $(n,T) = 1$ is inserted as a unit mask in the physical source before the graph pairing and every comparison.

The rest of the section proves this proposition. Denote the collection of fixed input constants by $\mathcal{C}_{\mathrm{in}}$. We normalize the reader bounds to one; the general bound follows by scaling.

7.1 The physical readers and the exact residual coefficient

Fix the depth J of the constructed graph and put $M = J + 1$. Extend each reader by zero when one of its count coordinates is negative. For positive x and negative y, define the physical readers

$$\begin{aligned} a^{[J]}(x) &= a\big((x^\circ)_{\le Q}, (\Omega_g(x) - J)_g\big),\\ b^{[J]}(y) &= b\big(|y|^\circ, (\Omega_g(|y|) - J)_g\big),\\ U_T(x) &= \mathbf{1}_{(x,T)=1}. \end{aligned} \quad (7.5)$$

These are functions of the integer position. They do not depend on which ordered label array represents a vertex above that position.

Choose $\delta \in (0, 1/2)$ with $\operatorname{supp}\chi \subseteq [\delta, 1-\delta]$, and put $C_\delta = 2(1-\delta)/\delta$. Let ψ be a real smooth function supported in $[1/2, 2]$, with $|\psi| \le 1$. At an even integer soft scale X, the source and target position functions used in the physical pairing are

$$\begin{aligned} \mathcal{F}_{X,\zeta}(x) &= F(x)a^{[J]}(x)U_T(x)\psi(x/X)(x/T)^{i\zeta} && (x > 0),\\ \mathcal{H}_X(y) &= \mathbf{1}_{[-C_\delta X, -1]}(y)b^{[J]}(y) && (y < 0), \end{aligned} \quad (7.6)$$

and they are zero elsewhere. The source vector is $d_\vartheta(x)(X/x)\mathcal{F}_{X,\zeta}(x)$, the target vector is $d_\vartheta(y)\mathcal{H}_X(y)$, and the pairing is divided by X. We use the physical step $-T$, additive parameter zero, and an inner product conjugate-linear in its first argument. These conventions fix the phase and the reciprocal factor used below.

For a comparison pattern $\nu = (A, B)$, let $C_{\nu,g}$ be the replaced slots and $S_{\nu,g}$ the shared slots, and put

$$c_g = |C_{\nu,g}|, \qquad s_g = |S_{\nu,g}| = J - c_g, \qquad t_g = 1 + c_g = M - s_g, \qquad \delta_{\nu,g} = s_g - J = -c_g.$$

The selected sets A and B are nonempty subsets of the two high blocks. All low groups are untouched. A shared tuple has product D; the free tuples have products a_f and b_f. The endpoints are exactly

$$x = Dn, \qquad y = D(n - Ta_f b_f). \quad (7.7)$$

Every prime in D, a_f, b_f avoids T. Consequently

$$U_T(Dn) = U_T(n), \qquad (n - Ta_f b_f, T) = (n, T).$$

The source mask therefore also restricts the target to units, in every branch. This remains true when T has many small prime factors.

Define the residual readers

$$\begin{aligned} a_\nu(n) &= a\big((n^\circ)_{\le Q}, (\Omega_g(n) + \delta_{\nu,g})_g\big), \\ b_\nu(n) &= b\big(n^\circ, (\Omega_g(n) + \delta_{\nu,g})_g\big). \end{aligned} \tag{7.8}$$

Since D contains exactly s_g labels in group g, counted with multiplicity, the following identities are exact even for repeated labels in an unrestricted shared tuple:

$$a^{[J]}(Dn) = a_\nu(n), \qquad b^{[J]}(-Dn) = b_\nu(n), \qquad F(Dn) = \gamma(D)F(n). \tag{7.9}$$

The identities use Ω_g, not ω_g. The residual readers retain no identities of the shared primes.

Lemma 7.2 (The covariant source after selecting one low prime)**.** *Fix a shared tuple and a complete source feature from a comparison branch. Choose an untouched low group g_0. After expansion of its first-order mark, write $n = upv$, where $p \in \mathcal{P}_{g_0}$ and u is supported on the prime divisors of an auxiliary modulus $r \le L^{D_r}$, with fixed D_r. On the retained unit terms, $(u, T) = 1$ and $(v, Tr) = 1$. The residual coefficient multiplying $F(v)$ can be taken to be*

$$R_{\nu,u}(v) = \vartheta^{\omega_{g_0}(v)} O_\rho(uv)\, a\big(((uv)^\circ)_{\le Q}, (\Omega_g(uv))_g + \mathbf{e}_{g_0} + \delta_\nu\big), \tag{7.10}$$

where O_ρ contains the complete marked sums of all the other groups. This is a function of $v^\flat$, of absolute value at most a fixed power of L, uniformly in the readers and the shared tuple. It is independent of the identity of p. The only source factors removed from it are the scalars $\gamma(D)$, $\gamma(u)$, and $\gamma(p)$; the last scalar is constant over $\mathcal{P}_{g_0}$. The replacement needed for distinct-prime damping is supported on $p^2 \mid n$ and has the square-energy bound proved below.

Proof. All prime divisors of u are at most $r < Q$, and $p \nmid rT$ for large L. Covariance gives $F(Dupv) = \gamma(D)\gamma(u)\gamma(p)F(v)$. Disjointness of the groups gives

$$(upv)^\circ = (uv)^\circ, \qquad \Omega_g(upv) = \Omega_g(uv) + \mathbf{1}_{g=g_0}.$$

The complete other-group feature does not read the selected low label, so $O_\rho(upv) = O_\rho(uv)$. All the quantities in (7.10) are determined by $v^\flat$ once u and the pattern have been fixed. In particular, the full exterior part of v is not observed by the source reader. The reader identity is still exact when $p \mid v$, because its counters include multiplicity. Only the identity $\omega_{g_0}(upv) - 1 = \omega_{g_0}(v)$ can then fail; this is the prime-square replacement. The polynomial bound for O_ρ follows from the complete endpoint bound below. □

7.2 Reindexing and complete endpoint features

The shared probability law is

$$\mu_{\mathrm{sh}}(\mathbf{s}) = \prod_g \prod_{i \in S_{\nu,g}} \frac{1}{V_g s_{g,i}} = \frac{\sigma_\nu}{D}, \qquad \sigma_\nu = \prod_g V_g^{-s_g}. \tag{7.11}$$

It has total mass one, including repeated labels. On the good-label part, shared labels are distinct and avoid both residual endpoints. Then $\omega_g(Dn) = s_g + \omega_g(n)$. The two damping factors at each endpoint combine to $\vartheta^{\omega_g(n) - t_g}$, and the vertex and edge normalizations satisfy

$$\prod_g V_g^{-M} V_g^{-t_g} = \sigma_\nu \prod_g V_g^{-2t_g}, \qquad \frac{1}{X}\sigma_\nu \frac{X}{Dn} = \frac{\mu_{\mathrm{sh}}(\mathbf{s})}{n}. \tag{7.12}$$

Thus the residual source measure is $1/n$.

For a nonzero integer n, let $\Lambda_\nu(n)$ be the finite set of lists choosing t_g ordered distinct primes dividing $|n|$ in each group. For any feature E on these lists, put

$$\mathcal{G}_E(n) = \sum_{\boldsymbol{\ell}\in\Lambda_\nu(n)} E(\boldsymbol{\ell}) \prod_g \frac{\vartheta^{\omega_g(n)-t_g}}{V_g^{t_g}}. \tag{7.13}$$

An empty list sum is zero. A cell kernel has the finite expansion

$$\begin{aligned} \mathcal{K}_C(d,v) &= \sum_{\eta\in\mathfrak{F}_C} u_{C,\eta}(d) v_{C,\eta}(v), \\ u_{C,j,\xi}(d) &= \mathbf{1}_{j(d)=j}\overline{\xi(d)}, \\ v_{C,j,\xi}(v) &= \mathbf{1}_{j(v)=j}\mathbf{1}_{\nu_{C,j}\ge\xi_{\mathrm{cell}}}\nu_{C,j}^{-1}\xi(v), \end{aligned} \tag{7.14}$$

where $\xi_{\mathrm{cell}} = L^{-x_*}$ is the cell-mass threshold and $\mathfrak{F}_C$ consists of occupied cells and primitive characters of conductor at most $Q_0 = \lceil L^{q_*}\rceil$. The cell threshold is distinct from the character ξ. Here $|u_{C,\eta}| \le 1$ and $|v_{C,\eta}| \le \xi_{\mathrm{cell}}^{-1}$. Each cell mass is computed with the actual coloured product law.

For source feature $\rho \in \mathfrak{F}_B$ and target feature $\eta \in \mathfrak{F}_A$, use the complete sums

$$G_{B,\rho}(n) = \mathcal{G}_{v_{B,\rho}\circ P_B}(n), \qquad G_{A,\eta}(r) = \mathcal{G}_{v_{A,\eta}\circ P_A}(r).$$

Fix D and put $Z = X/D$. The two sequences for the comparison are

$$\begin{aligned} f_{\rho,D}(n) &= \gamma(D)F(n)a_\nu(n)U_T(n)\frac{\psi(n/Z)}{n}(Dn/T)^{-i\zeta}G_{B,\rho}(n) && (n>0), \\ g_{\eta,D}(r) &= G_{A,\eta}(r)\mathbf{1}_{[-C_\delta X,-1]}(Dr)b_\nu(|r|) && (r<0), \end{aligned} \tag{7.15}$$

and zero otherwise. Every residual endpoint list is summed in these sequences before any square norm or Cauchy–Schwarz inequality is taken. In particular, no maximum over individual endpoint labels replaces a complete endpoint sum.

Let μ_A, μ_B be the original free-tuple probability laws. Direct finite reindexing gives

$$\begin{aligned} \mathcal{C}_{D,\eta,\rho} = \sum_{\alpha,\beta} \mu_A(\alpha)\mu_B(\beta) a_f^{i\zeta} u_{A,\eta}(a_f) b_f^{i\zeta} u_{B,\rho}(b_f) \\ \cdot \sum_n f_{\rho,D}(n) g_{\eta,D}(n - Ta_fb_f), \end{aligned} \tag{7.16}$$

and the unrestricted normalized branch is

$$\mathcal{U}_{\nu,X}(\zeta) = -(-1)^{|A|+|B|}\sum_{\mathbf{s}} \mu_{\mathrm{sh}}(\mathbf{s}) D^{i\zeta} \sum_{\eta,\rho} \mathcal{C}_{D,\eta,\rho}. \tag{7.17}$$

This is the coefficient dictionary of the two cell kernels, including which kernel reads which endpoint. The position functions (7.6) are unchanged by the slot permutations at either endpoint, so the same identity holds after those averages. These are the finite identities of [7, (7.11)–(7.15)], with the transported position readers and the unit mask kept inside the endpoints.

We record why removing the good-label restrictions has negligible cost. Put $P = \exp(L^{.27})$ and $D_{\max} = \exp(O_J(L^{.46}\log L))$. For fixed shared and free tuples, a forbidden incidence at a residual endpoint imposes one congruence modulo $p \ge P$ in an interval of length $O_\delta(X/D)$. It has $O_\delta(X/(Dp)+1)$ solutions. The reciprocal source and the fixed label normalizations then

give a normalized cost bounded by a fixed power of L times $P^{-1} + D_{\max}/X$. A coincidence of two sampled labels in a group has probability at most

$$\sum_{p\in\mathcal{P}_g} \mu_g(p)^2 \le \frac{1}{PV_g}.$$

For an incidence between a shared prime and a residual endpoint, the squared endpoint mass is at most a fixed multiple of $X\Delta_\nu$, where

$$\Delta_\nu = \sum_g \frac{s_g}{V_g} \sum_{p\in\mathcal{P}_g} p^{-2} \le \frac{\sum_g s_g}{P}.$$

Cauchy–Schwarz with the physical row bound gives $O(L^{C_J}\sqrt{\Delta_\nu})$. These counts are the finite restriction estimates of [7, (7.7)–(7.10)]; multiplying by the position readers or by U_T does not increase any absolute mass. Since J is fixed and $X \ge e^{9L}$, they give, uniformly in ζ,

$$|\mathcal{R}_{\nu,X}(\zeta) - \mathcal{U}_{\nu,X}(\zeta)| \ll_{J,\delta} L^{C_J}\left(e^{-L^{.27}/2} + D_{\max}/X\right) \ll_{A,J,\delta} L^{-A} \tag{7.18}$$

for every fixed $A > 0$. This also applies to the raw branch.

Here are the complete square masses. Let $W_\nu(n) = \mathcal{G}_1(n)$ and $C_m = (m+1)!4^{m+1}$, where $t_g \le m+1$. The elementary inequality

$$(w)_t\vartheta^{w-t} = t!\binom{w}{t}\vartheta^{w-t} \le \frac{t!}{(1-\vartheta)^t}$$

gives $W_\nu(n) \le C_m^s$. Conversely, summing first over a fixed distinct label list with product d and then over its multiples gives

$$\sum_{0<|n|\le H} W_\nu(n) \le 2H, \qquad \sum_{0<|n|\le H} |\mathcal{G}_E(n)|^2 \le 2H\xi_{\mathrm{cell}}^{-2}C_m^s \quad (|E| \le \xi_{\mathrm{cell}}^{-1}). \tag{7.19}$$

Indeed, the damping is at most one and the list reciprocal sum is at most $\prod_g V_g^{t_g}$, which cancels its normalization. On the source $1/n \le 2/Z$. Hence there is a fixed B_2, determined after m, x_*, for which

$$\|f_{\rho,D}\|_2^2 \ll_{C_F} Z^{-1}L^{B_2}, \qquad \|g_{\eta,D}\|_2^2 \ll_\delta ZL^{B_2}. \tag{7.20}$$

For example, any $B_2 > 2x_* + 2\log C_m + 2$ suffices. Factoring the untouched low group out of the complete source gives

$$G_{B,\rho}(n) = w_0(n)O_\rho(n), \qquad w_0(n) = \frac{\omega_{g_0}(n)\vartheta^{\omega_{g_0}(n)-1}}{V_{g_0}}, \tag{7.21}$$

where $w_0(n) = 0$ if the count is zero. The complete other-group sum O_ρ is bounded by a fixed power of L and depends only on $n^\flat$. This proves the bound used in Lemma 7.2.

7.3 The exact unit energy at an arbitrary total

We establish the arithmetic estimate for the complete source of a comparison branch $\nu \ne 0$, on a dyad of free products $a_f \in [U, 2U)$, $b_f \in [V, 2V)$ in (7.16). The support argument at the start of Section 7.4 shows that only dyads satisfying (7.36) can contribute. Discard the other dyads before applying the present estimates. The height bounds below are asserted only on

retained dyads and are verified in that section. The raw branch is recovered separately from the full moment in Section 7.5, after the comparison estimates.

For an auxiliary modulus $r \le L^{D_r}$, retain only integers with

$$u = u_r(n) = \prod_{p \mid r} p^{v_p(n)} \le U_0, \qquad U_0 = \exp(L^{.39}/4).$$

The source already contains U_T. Thus only $(u, T) = 1$ occurs, and after selecting p the remaining factor v satisfies $(v, Tr) = 1$. Set

$$d = (u, r), \qquad q = Tr/d, \qquad dq = Tr. \tag{7.22}$$

Then $(u/d, q) = 1$. Indeed, comparison of valuations gives $(u/d, r/d) = 1$, while $(u, T) = 1$. There is no assumption that T and r are coprime.

Partition $\mathcal{P}_{g_0}$ into clipped dyadic slices $\mathcal{P}_j \subseteq (a_j, b_j]$, with $b_j \le 2a_j$, and use its complete mass $V_0 = V_{g_0}$ to normalize every slice. Set $H_0 = \lfloor 8Z \rfloor$ and

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

This common residual set includes all nonzero terms with $\psi(upv/Z) \ne 0$. On every retained comparison dyad, we will check uniformly in the selected prime and the retained supported part that

$$p^2 uTr \le H_0, \qquad e^{3L/4} \le Z/(up) \le e^{5L/4}. \tag{7.23}$$

In particular, $q, \ell_j \le 8Z/(ua_j)$.

Up to a common scalar κ_u of absolute value one for each u, the character slices are

$$S_{u,j}(\xi) = \frac{1}{V_0} \sum_{p \in \mathcal{P}_j} \sum_{v \in \mathcal{V}_j} F(v) R_{\nu,u}(v) \frac{\psi(upv/Z)}{upv} e(\beta upv)(upv)^{i\tau} \xi(p)\xi(v). \tag{7.24}$$

Here $|\beta|Z$ and $|\tau|$ have fixed polynomial bounds. The corresponding additive polynomial on $\mathbb{Z}/(Tr)\mathbb{Z}$ is obtained by replacing $\xi(p)\xi(v)$ by $e((u/d)pvh/q)$, and is denoted $A_u(h)$. Collect the coefficients of *all* slices with the same product residue $pv \pmod q$ into a function on $(\mathbb{Z}/q\mathbb{Z})^\times$. Additive orthogonality and unit-group character orthogonality give the exact identity

$$\boxed{\sum_{h \bmod Tr} |A_u(h)|^2 = \frac{dq}{\varphi(q)} \sum_{\xi \bmod q} \left| \sum_j S_{u,j}(\xi) \right|^2.} \tag{7.25}$$

To see both normalizations, write $c(x)$ for the sum of all the coefficients on the product class $x = pv \bmod q$, including all slices. Additive orthogonality gives $dq \sum_x |c(x)|^2$, since u/d is a unit and each frequency modulo q occurs d times. Multiplicative orthogonality gives $\varphi(q) \sum_x |c(x)|^2$ for the character sum on the right before its prefactor.

For example, take $T = 6$, $r = 15$, $u = 5$. Then $d = 5$, $q = 18$ and the prefactor is 15. Two slice coefficients equal to 1 at class 1, and one equal to -1 at class 5, have additive energy $90(2^2 + 1) = 450$. Squaring the slices separately would give 270 if the two positive terms belong to different slices. This example concerns the exact finite identity, independently of the subsequent prime estimates. We denote the principal character by ξ_0. For $q = 1$, take $q/\varphi(q) = 1$ and an empty nonprincipal sum. Thus (7.25) is precisely the unit-group identity of [7, (6.11)], now justified without a roughness assumption on the total.

The factor $q/\varphi(q)$ cannot be discarded for an arbitrary total. Uniformly for $q \le e^{2L}$,

$$\frac{q}{\varphi(q)} \ll \log L. \tag{7.26}$$

For example, primes at most L^2 contribute $O(\log L)$ by the Mertens product bound. There are at most $2L/\log(L^2)$ larger prime divisors, whose reciprocal sum is $O(1/(L\log L))$, so their product is bounded. The weaker bound $O(L)$ would also suffice.

To apply (7.3) in the principal character, absorb the restriction from prime divisors of Tr at most Q into the coefficient of $v^\flat$. All prime divisors of r are in this range. The remaining unit restriction has absolute error at most a fixed polynomial in L times

$$Z_p \sum_{\substack{p' \mid T \\ p' > Q}} \frac{1}{p'} \ll Z_p L/Q, \qquad Z_p = Z/(up). \tag{7.27}$$

The bound follows by counting multiples up to a fixed multiple of Z_p; no $O(1)$ error per residue is needed. After increasing the precision in (7.3), the test $\widetilde{\psi}(v/Z_p)e(\beta upv)v^{i\tau}$, where $\widetilde{\psi}(t) = \psi(t)/t$, gives

$$|S_{u,j}(\xi_0)| \ll_{A_0} u^{-1}L^{-A_0}V_0^{-1}\sum_{p\in\mathcal{P}_j} p^{-1}.$$

Its amplitude and total variation have fixed polynomial bounds. Summing the complete slices and using their total reciprocal mass V_0 gives

$$\frac{dq}{\varphi(q)}\left|\sum_j S_{u,j}(\xi_0)\right|^2 \ll_{A_0} \frac{d}{u^2}(\log L)L^{-2A_0}. \tag{7.28}$$

Thus the totient loss is absorbed by choosing A_0 after the desired local energy precision. If an additional finite slice decomposition is retained, its polynomial count is absorbed in the same choice.

For completeness, the nonprincipal estimate keeps the same quotient factor throughout. For $w = \beta Z$, put

$$\Psi_w(t) = \widetilde{\psi}(e^t)e(we^t).$$

Fourier inversion in the logarithmic coordinate gives

$$S_{u,j}(\xi) = \frac{1}{Z}\int_{\mathbb{R}} \widehat{\Psi}_{\beta Z}(t)Z^{-2\pi it}u^{i(\tau+2\pi t)}\left(\frac{1}{V_0}\sum_{p\in\mathcal{P}_j} p^{i(\tau+2\pi t)}\xi(p)\right) \cdot\left(\sum_{v\in\mathcal{V}_j} F(v)R_{\nu,u}(v)v^{i(\tau+2\pi t)}\xi(v)\right)dt. \tag{7.29}$$

The factor $1/Z$ is forced by the reciprocal source. For coefficients bounded by L^{C_0} on an interval of length ℓ_j, with zero coefficients off any chosen subset of the units, character orthogonality gives

$$\sum_{\xi \bmod q}\left|\sum_v b_v\xi(v)\right|^2 \le \varphi(q)(1+\ell_j/q)\ell_j L^{2C_0}. \tag{7.30}$$

Each residue class contains at most $1+\ell_j/q$ interval points. Multiplying by $dq/\varphi(q)$ cancels the totient denominator in this estimate. The height condition gives

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

Consequently (7.1), (7.29), and (7.30) give central-frequency energy

$$\ll \frac{d}{u^2} L^{C_2} \exp(-2cL^{.27}).$$

The factors a_j^2 and Z^2 cancel against the residual interval length and the external reciprocal normalization.

There are at most a fixed polynomial number of slices. If $|\beta|Z, |\tau| \le L^{B_0}$, truncate the integral at $|t| \le L^{D_0}$ with a fixed $D_0 > B_0 + 1$. For every integer $k \ge 2$,

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

On the tail use the trivial prime bound $2a_j$ and the same unit orthogonality estimate. Choose k after the requested accuracy. Minkowski in the finite character space, followed by Cauchy–Schwarz over the finite slices, now proves

$$\frac{dq}{\varphi(q)} \sum_{\xi \ne \xi_0} \left| \sum_j S_{u,j}(\xi) \right|^2 \ll_A \frac{d}{u^2} L^{-A}. \tag{7.31}$$

All truncated frequencies lie in the range of (7.1) for large L. Combining this with (7.28) in the exact identity proves

$$\sum_{h \bmod Tr} |A_u(h)|^2 \ll_A \frac{(u,r)}{u^2} L^{-A}. \tag{7.32}$$

Summing over supported parts uses Minkowski, not a sum of their individual squared norms. The convergent Euler product satisfies

$$\sum_{\substack{u \ge 1 \\ p|u \Rightarrow p|r}} \frac{\sqrt{(u,r)}}{u} = \prod_{p^a \| r} \left(\sum_{k=0}^{a} p^{-k/2} + p^{a/2} \sum_{k>a} p^{-k} \right) \le r^2.$$

For $p \ge 3$ the displayed local factor is at most p^{2a} by a geometric sum; for $p = 2, a = 1$ it is $1 + \sqrt{2} < 4$, and for $a \ge 2$ it is at most $(1 - 2^{-1/2})^{-1} < 16 \le 2^{2a}$. The restriction $(u, T) = 1$ only decreases this positive sum. Thus

$$\sum_{h \bmod Tr} \left| \sum_{\substack{u \le U_0 \\ (u,T)=1}} \kappa_u A_u(h) \right|^2 \ll_A r^4 L^{-A}. \tag{7.33}$$

We check the two replacements in the complete physical coefficient. Its absolute value is $O(L^{C_3}/Z)$. Parseval and Rankin's inequality give

$$\int_{\mathbb{R}/\mathbb{Z}} |\widehat{f}_{\mathrm{large,r}}(\alpha)|^2 d\alpha \ll Z^{-1} L^{2C_3} U_0^{-1/2} r^2 \ll_A Z^{-1} L^{-A}, \tag{7.34}$$

because $\#\{n \le H_0 : u_r(n) > U_0\} \le H_0 U_0^{-1/2} \prod_{p|r} (1 - p^{-1/2})^{-1} \le H_0 U_0^{-1/2} r^2$. There is no deleted nonunit coefficient: those terms were zero in the original physical source.

For the prime-square replacement in Lemma 7.2, write the selected-prime average with probabilities $1/(V_0 p)$. For fixed p, its error coefficient is $O(pL^{C_3}/Z)$ and supported on multiples of p^2. Since $(p, Tr) = 1$, each residue modulo Tr contains at most $1 + H_0/(p^2 Tr) \le 2H_0/(p^2 Tr)$

such multiples, by (7.23). Cauchy–Schwarz in each residue and additive orthogonality give residue energy $O(L^{2C_3}/p^2)$. Jensen for the selected-prime probabilities therefore gives

$$\sum_{h \bmod Tr} |A_{\mathrm{square}}(h)|^2 \ll L^{2C_3} P^{-2} \ll_A L^{-A}. \tag{7.35}$$

Both masks, U_T and $u_r(n) \le U_0$, stay inside this calculation. The reader identity in the square error is exact; no invariance of a reader under a prime square is asserted.

7.4 Major and minor arcs for one complete comparison

Restrict the two free products in (7.16) to dyads $a_f \in [U, 2U)$ and $b_f \in [V, 2V)$. A nonzero summand has $n \in (Z/2, 2Z)$ and $n - Ta_f b_f \in [-C_\delta Z, 0)$, whence

$$Z/8 < TUV < (2 + C_\delta)Z. \tag{7.36}$$

Every dyad outside this range has zero contribution to (7.16) and is discarded before the energy estimates of Section 7.3 are applied. The selected sets are nonempty in the high groups, so $U, V \ge \frac{1}{2} e^{L^{.39}}$. Their logarithms are $O_m(L^{.46} \log L)$, and therefore

$$1 + \log(2UV) \le L^{d_f}, \qquad d_f = .48$$

for large L. Unique factorization identifies the multiset of factors in each group. At most $m!$ orders in a group produce the same value, giving, for a fixed C_a determined by m,

$$\sum_{\alpha : a(\alpha) = a_f} \mu_A(\alpha) \le L^{C_a}/a_f, \qquad \sum_{\beta : b(\beta) = b_f} \mu_B(\beta) \le L^{C_a}/b_f. \tag{7.37}$$

This includes repeated labels in the free laws.

Given a desired minor-arc precision $B' > 0$, put

$$\Gamma = 2C_a + 2B' + 8, \quad R_0 = \lceil L^\Gamma \rceil, \quad S = L^{d_f + \Gamma}, \quad \Delta_r = \frac{2S}{TrUV} \quad (1 \le r \le R_0). \tag{7.38}$$

The harmless factor two avoids integer-rounding conventions in Dirichlet approximation. Let $\mathfrak{M}$ be the union of the arcs of radius Δ_r about all rationals with denominator Tr, $r \le R_0$. The free-label multiplier on this dyad is

$$\mathcal{M}(\alpha) = \sum_{\alpha_f, \beta_f} \mu_A(\alpha_f) \mu_B(\beta_f) a_f^{i\zeta} u_{A,\eta}(a_f) b_f^{i\zeta} u_{B,\rho}(b_f) e(-T a_f b_f \alpha). \tag{7.39}$$

It satisfies $|\mathcal{M}| \le 1$ everywhere.

We use the elementary bilinear bound: for product distributions of mass at most one on $[U, 2U)$ and $[V, 2V)$ with atoms at most K/U and K/V, and for a reduced approximation $|\alpha - a_0/r| \le r^{-2}$,

$$\left| \sum_{a,b} c_a d_b e(\alpha a b) \right| \ll K \left(\frac{1}{V} + \frac{1}{r} + \frac{1 + \log(2r)}{U} + \frac{r(1 + \log(2r))}{UV} \right)^{1/2}. \tag{7.40}$$

Cauchy–Schwarz in a reduces the square to a sum over differences $b - b'$. The geometric sum in a is bounded by $\min(2U, \|\alpha(b - b')\|^{-1})$. Sorting the differences into blocks modulo r and summing reciprocal distances gives the four terms in (7.40); the atom bounds supply K^2 before taking the square root. Applying this after grouping the original tuples by product proves it for the free laws here.

Dirichlet approximation to $T\alpha$ with denominator at most UV/S puts α in $\mathfrak{M}$ whenever its approximating denominator is at most R_0. On the complement the denominator is larger than R_0 and at most UV/S. With $K = L^{C_a}$, the terms $1/U, 1/V$ have arbitrary logarithmic saving and the other terms are $O(L^{-\Gamma})$. Hence

$$\sup_{\alpha\notin\mathfrak{M}} |\mathcal{M}(\alpha)| \le L^{-B'}. \tag{7.41}$$

The height checks needed in the preceding energy calculation now follow from this geometry. For the selected low prime, a retained supported part, and $r \le R_0$,

$$\log \frac{p^2 ur}{UV} \le 2L^{.36} - \tfrac{7}{4}L^{.39} + O_\Gamma(\log L) \longrightarrow -\infty.$$

Together with the upper bound in (7.36), this gives $p^2uTr = o_\delta(Z)$ uniformly on retained dyads, and hence $p^2uTr \le 4Z \le H_0$ for large L. For physical scales $e^{.9L} \le X \le e^{1.1L}$, the bound $\log D = O_J(L^{.46}\log L)$ also gives $Z/(up) \in [e^{3L/4}, e^{5L/4}]$. Thus (7.23) holds uniformly on the retained comparison dyads. The lower bound in (7.36) gives $Z\Delta_r \le 16S/r$.

For fixed r, integrate the complete residue-energy bound, including (7.35), over $|\beta| \le \Delta_r$. Use (7.34) on the whole circle for the discarded supported parts. If $E_r \ll_A r^4 L^{-A}$ denotes the retained residue bound, the union bound and the two-term square inequality give

$$\int_{\mathfrak{M}} |\widehat{f}_{\rho,D}(\alpha)|^2 d\alpha \ll \sum_{r\le R_0} \Delta_r E_r + \sum_{r\le R_0} \int_{\mathbb{R}/\mathbb{Z}} |\widehat{f}_{\mathrm{large,r}}|^2.$$

The deleted term is bounded once on the union for each r, rather than once for every rational numerator. Since R_0, S, and the supported-part losses have fixed polynomial bounds, increasing the local accuracy gives, for every fixed $A_1 > 0$,

$$\int_{\mathfrak{M}} |\widehat{f}_{\rho,D}(\alpha)|^2 d\alpha \ll_{A_1} Z^{-1} L^{-A_1}. \tag{7.42}$$

This includes the factor $q/\varphi(q)$ for arbitrary T and the complete source-reader coefficient.

With $\widehat{f}(\alpha) = \sum_n f(n)e(n\alpha)$, finite Fourier orthogonality gives the exact comparison identity

$$\mathcal{C}^{U,V}_{D,\eta,\rho} = \int_{\mathbb{R}/\mathbb{Z}} \widehat{f}_{\rho,D}(\alpha)\widehat{g}_{\eta,D}(-\alpha)\mathcal{M}(\alpha)\, d\alpha. \tag{7.43}$$

Cauchy–Schwarz separately on the two arc sets, followed by Parseval, (7.20), (7.41), and (7.42), yields the explicit estimate

$$\boxed{|\mathcal{C}^{U,V}_{D,\eta,\rho}| \ll_{C_F,\delta} L^{(B_2-A_1)/2} + L^{B_2-B'}.} \tag{7.44}$$

The physical length Z cancels on both arc sets. This is the comparison estimate in [7, (7.27)–(7.30)], with the complete marked source energy and its unit mask specified above.

Let L^{C_Σ} bound the product of the number of comparison patterns, the two feature counts, and the two free-product dyad counts. This exponent is fixed once m, q_*, h_*, x_* are fixed: there are at most $2^{m|\mathcal{B}|}$ patterns, at most $L^{h_*+2q_*+.46+4}$ features for either selected product, and at most $L^{1.46}$ occupied dyads for each product, eventually. Given a desired total comparison precision A_c, choose

$$B' \ge B_2 + C_\Sigma + A_c + 3, \qquad A_1 \ge B_2 + 2(C_\Sigma + A_c + 3). \tag{7.45}$$

Choose the residue and principal precisions after these two exponents. Then sum (7.44) over all features, patterns, and retained product dyads, and finally average over the shared law of mass one. Formula (7.17) and (7.18) give

$$\sum_{\nu\neq 0} |\mathcal{R}_{\nu,X}(\zeta)| \ll_{A_c,\delta} L^{-A_c} \qquad (|\zeta| \le L^2). \tag{7.46}$$

The choices in (7.45) change only arithmetic accuracies and the final threshold. They do not change the kernels, the depth, or the graph moment required of the construction.

7.5 Using the physical moment

We use the constructed form of Proposition 6.1, at a precision E to be fixed below. Its pairing consequence [7, (7.32)] is recalled to specify the position-reader requirement. In a weighted graph whose edges have displacements kd with $k \neq 0$ and $|d| \le \mathcal{B}$, let u_z denote the indicator of the position fibre and $T_R = \sum_{z\in I}\langle u_z, (AA^*)^R u_z\rangle$. For $f = \sum_{z\in I} a_z u_z$, with $|a_z| \le 1$, one has

$$|\langle f, Ag\rangle| \le \sqrt{3}(\mathcal{B}+2)^{1/(2R)}\|g\|\,\|f\|^{1-1/R} T_R^{1/(2R)}. \tag{7.47}$$

Indeed, split each residue class modulo $|k|$ into consecutive blocks of $b = \max(1, \lceil \mathcal{B}\rceil)$ positions in its own lattice. Write B_j for the jth block, f_j for the restriction of f to B_j and g_j for the restriction of g to that block and its two neighbours. Only g_j can be paired with f_j. For $P = AA^*$, the spectral theorem gives

$$\|A^* f_j\| \le \|f_j\|^{1-1/R}\|P^{R/2} f_j\|^{1/R}, \qquad \|P^{R/2} f_j\|^2 \le b \sum_{z\in B_j\cap I} \langle u_z, P^R u_z\rangle.$$

The second inequality is Cauchy–Schwarz applied to $\sum a_z P^{R/2} u_z$. Summing $\|g_j\|\|A^* f_j\|$ uses Hölder with exponents $2, 2R/(R-1), 2R$ (the middle factor is omitted for $R = 1$). Now $\sum_j \|f_j\|^2 = \|f\|^2$ and $\sum_j \|g_j\|^2 \le 3\|g\|^2$. This proves (7.47). This argument requires a position coefficient a_z, rather than an arbitrary label-dependent reader. Our readers in (7.5) have exactly that form.

For a hard source interval $[X_0, 2X_0)$, the coefficient $d_\vartheta(x)(X_0/x)\mathcal{F}_{X,\zeta}(x)$ is bounded by C_F. Normalize this constant first. For a list of distinct labels of product D, the number of eligible positive positions up to $2X_0$ is at most $2X_0/D$. Summing over lists with vertex mass τ costs at most one in reciprocal mass, since the corresponding product laws have mass one. The damping is at most one. The squared norm of the normalized source is therefore at most $2X_0$ and the same count gives at most $C_\delta X \le 2C_\delta X_0$ for the common target. The moment bound $T_R \le X_0 L^{-2ER}$ therefore gives, after division by X_0,

$$|\mathcal{R}_{\mathrm{full},X_0}(\zeta)| \ll_{C_F,\delta} (\mathcal{B}+2)^{1/(2R)} L^{-E}.$$

Here $R = \lceil L^{.52}\rceil$ and $\log \mathcal{B} = O_J(L^{.46}\log L)$, so the extra factor tends to one. The hard source intervals $[X/2, X)$ and $[X, 2X)$ both use the target $[-C_\delta X, -1]$. Adding them gives the soft source pairing because $X_0^{-1}(X_0/x) = 1/x$ at both scales. Hence

$$|\mathcal{R}_{\mathrm{full},X}(\zeta)| \ll_{C_F,\delta} L^{-E} \qquad (|\zeta| \le L^2). \tag{7.48}$$

Here is a noncircular order of choices for a marked accuracy $A > 0$. First fix $E = A + 6$. Choose the ideal kernel parameters m, q_*, h_*, x_* at the precision needed for the physical moment, and then fix an admissible J. Now B_2, C_a, C_Σ are fixed. Set $A_c = A + 6$ in (7.45); this fixes B', A_1, the major-arc exponents, and then the local residue and principal accuracies. Finally choose Fourier-tail orders and the threshold in L. No later accuracy changes a previously

chosen graph parameter. For each admitted prime family, the cell coefficients are formed before the frequency, the position readers, and their pairing are chosen; the same family of graph coefficients works for all those readers. Since the full graph is the raw branch plus all comparison branches, (7.18), (7.46), and (7.48) give

$$|\mathcal{U}_{0,X}(-2\pi t)| \ll_{A,C_F,\delta} L^{-A-6} \qquad (|t| \le L). \tag{7.49}$$

The range $2\pi L \le L^2$ holds for large L.

7.6 The raw weight, the shared sign, and counting measure

In the raw branch $s_g = J$ and $t_g = 1$ in every group. Thus each complete residual list sum in (7.13) is exactly

$$W(n) = \prod_g \frac{\omega_g(n)\vartheta^{\omega_g(n)-1}}{V_g}, \tag{7.50}$$

with a zero factor when the count is zero. There are no free labels, and the residual reader shifts are zero. Moreover

$$\gamma(D) = \Gamma_0, \qquad \Gamma_0 = 1 \quad (\gamma = 1), \qquad \Gamma_0 = (-1)^{Js} \quad (\gamma = \lambda). \tag{7.51}$$

This scalar is independent of the shared tuple, including tuples with repetitions. It has absolute value one.

The kernel phase $D^{-2\pi it}$ cancels the D in the conjugated source phase $(Dn/T)^{2\pi it}$. The unrestricted raw identity is therefore

$$\begin{aligned}\mathcal{U}_{0,X}(-2\pi t) = \Gamma_0 \sum_{\mathbf{s}} \mu_{\mathrm{sh}}(\mathbf{s}) \sum_{0<n<T} \frac{W(n)W(T-n)}{n} \\ \cdot F(n)a(Y_{\mathrm{src}}(n))b(Y_{\mathrm{tgt}}(T-n))U_T(n) \\ \cdot \psi(Dn/X)\mathbf{1}_{[-C_\delta X,-1]}(D(n-T))e\big(t\log(n/T)\big).\end{aligned} \tag{7.52}$$

In particular, the shared factor does not survive inside the source reader or the phase.

The Mellin profile that recovers counting measure is

$$\Phi(t) = e^t\chi(e^t), \qquad \widehat{\Phi}(s) = \int_{\mathbb{R}} \Phi(t)e(-st)\,dt. \tag{7.53}$$

It is a fixed smooth function of compact support. Fourier inversion gives the exact normalization

$$\boxed{\frac{1}{n}\int_{\mathbb{R}} \widehat{\Phi}(s)e\big(s\log(n/T)\big)\,ds = \frac{1}{T}\chi(n/T).} \tag{7.54}$$

The profile $\chi(e^t)$ alone would retain a factor $1/n$. The choice (7.53) agrees with [7, (7.38)]. On the support of $\chi(n/T)$, $n/T \in [\delta, 1-\delta]$; if $\psi(Dn/X) \ne 0$, then $D(T-n) \le (1-\delta)Dn/\delta < C_\delta X$. The common-target indicator in (7.52) is thus one after inversion. Integrating that identity yields

$$\begin{aligned}\int_{\mathbb{R}} \widehat{\Phi}(t)\mathcal{U}_{0,X}(-2\pi t)\,dt = \frac{\Gamma_0}{T} \sum_{\substack{0<n<T \\ (n,T)=1}} \chi(n/T)W(n)W(T-n)F(n)a(Y_{\mathrm{src}}(n))b(Y_{\mathrm{tgt}}(T-n)) \\ \cdot \sum_{\mathbf{s}} \mu_{\mathrm{sh}}(\mathbf{s})\psi(Dn/X).\end{aligned} \tag{7.55}$$

All sums before the integral are finite.

For the complementary frequencies, $W(n) \le 4^s$. For fixed D, the soft source has $Z/2 < n < 2Z$ and $1/n \le 2/Z$. The shared law has mass one. Consequently

$$|\mathcal{U}_{0,X}(-2\pi t)| \ll C_F 4^{2s} \le L^{C_4} \qquad (t \in \mathbb{R})$$

for a fixed C_4. The rapid Fourier decay of the fixed Φ makes the integral over $|t| > L$ smaller than every prescribed negative power of L. Together with (7.49), this bounds the absolute value of (7.55) by $O_{A,\chi}(L^{-A-5})$.

Finally take a smooth dyadic partition $\psi(v) = \eta(v) - \eta(2v)$, where $\eta = 1$ on $(-\infty, 1]$, $\eta = 0$ on $[2, \infty)$, and $0 \le \eta \le 1$. For $K = 1 + \lceil 2L/\log 2 \rceil$ and $X_j = 2^j$,

$$\sum_{j=1}^{K} \psi(x/X_j) = 1 \qquad (2 \le x \le e^{2L}). \tag{7.56}$$

On the support in (7.55), every product Dn lies in this range for large L. If its jth cutoff is nonzero, then

$$\delta T/2 < X_j < 2(1-\delta)TD_{\max}.$$

Since $\log T = L + O(L^{\cdot 20})$ and $\log D_{\max} = o(L)$, all these physical scales are in $[e^{\cdot 9L}, e^{1.1L}]$, and the two hard source scales are in the range of the moment proposition. All other soft-scale integrals vanish after inversion.

Apply (7.56) *for each fixed shared tuple* before averaging over its law. The partition sum is then one, and only then does $\sum_{\mathbf{s}} \mu_{\mathrm{sh}}(\mathbf{s}) = 1$ remove the shared law. The scalar Γ_0 in (7.51) may be removed by absolute values. There are $O(L)$ scales, so we have proved

$$\left| \frac{1}{T} \sum_{\substack{0<n<T \\ (n,T)=1}} \chi(n/T) W(n) W(T-n) F(n) a(Y_{\mathrm{src}}(n)) b(Y_{\mathrm{tgt}}(T-n)) \right| \ll_{A,\chi,C_F,\mathcal{C}_{\mathrm{in}}} L^{-A}. \tag{7.57}$$

All bounds used the readers only through their displayed information and their supremum norms. Their functions, the total, and the regular colouring may therefore vary after the common threshold. This proves Proposition 7.1. For the three sources F_λ, F_0, F_* the uniform input constants depend only on the fixed smoothness parameter ε. Applying the proposition with $A = 1$ gives the marked precision required in Proposition 2.2.

The position-source lemma is essential to this conclusion. An arbitrary coefficient of the original $n^\flat$ need not be independent of the selected prime, whereas the coefficient in (7.10) is. The estimate concerns the product of two separately admissible readers; it does not cover arbitrary joint functions of their observations.

8 Recovering the joint laws and their local variants

The passage from weighted correlations to ordinary counting must preserve a little more information than the signs themselves. Squarefreeness, finitely many divisibility conditions, and the changes of prime values allowed below are all determined, up to a small exceptional population, by a sufficiently short prime part. Throughout this section the two-site conclusions assume the corresponding instances of (2.11). The squarefree density in Lemma 8.2 and the local reduction in Lemma 8.4 are independent of this premise. The functions of the two prime parts are separate; no estimate against an arbitrary function of both parts is used.

Fix $0 < \varepsilon < 1/4$, put $L = \log N$, and write

$$B = \exp(L^{0.26}), \qquad y = N^{\varepsilon}, \qquad n_{\le B} = \prod_{p \le B} p^{v_p(n)}, \qquad S(n) = \mathbf{1}_{P^+(n) \le y}, \qquad r = \rho(1/\varepsilon).$$

We take $P^+(1) = 1$. Both B and y use the outer total N throughout this section, including when the argument of S is $N - n$. The exponent 0.26 places B below the two prime ranges used in the graph construction. For a fixed real $\chi \in C_c^\infty((0,1))$, set

$$\mathcal{E}_{N,\chi}(H) = \frac{1}{N} \sum_{1 \le n < N} \chi(n/N) H(n, N-n).$$

The local estimate required below is

$$\sup_{|a|,|b| \le 1} \left| \mathcal{E}_{N,\chi}(F(n) a(n_{\le B}) b(m_{\le B}) V(m)) \right| = o_{\varepsilon,\chi}(1), \qquad m = N - n, \tag{8.1}$$

for each

$$F \in \{S - r, \lambda, S\lambda\}, \qquad V \in \{1, S, \lambda, S\lambda\}.$$

Here a and b are real functions on the positive integers, may depend on N, and each is a function only of its displayed argument. The error is uniform over these choices. The same assertion with the sites interchanged follows by replacing $\chi(t)$ by $\chi(1-t)$. Under the marked premise, Proposition 4.2 proves (8.1) after removal of the marking weights and replacement of the model centres. In particular, it is not inferred merely from the unmodulated assertion $C_\lambda(N) = o(N)$.

8.1 A finite basis for the sixteen states

There are four states at one site: a smoothness value in $\{0,1\}$ and a Liouville sign. The functions

$$1, \quad S - r, \quad \lambda, \quad S\lambda$$

form a basis for functions on this four-point state space. This small algebra is the only reconstruction step needed for the complete law.

Lemma 8.1 (Local masks and state probabilities)**.** *Let $A(n) = a(n_{\le B})$ and $D(m) = b(m_{\le B})$, where $|a|, |b| \le 1$. Put*

$$I_{i,\sigma}(v) = \mathbf{1}_{S(v)=i,\, \lambda(v)=\sigma}, \qquad \pi_1 = r, \quad \pi_0 = 1 - r.$$

For $i, j \in \{0,1\}$ and $\sigma, \tau \in \{-1, 1\}$,

$$\mathcal{E}_{N,\chi}(A(n) D(m) I_{i,\sigma}(n) I_{j,\tau}(m)) = \frac{\pi_i \pi_j}{4} \, \mathcal{E}_{N,\chi}(A(n) D(m)) + o_{\varepsilon,\chi}(1). \tag{8.2}$$

The error is uniform over the two masks.

Proof. The identities

$$I_{1,\sigma} = \frac{r}{2} + \frac{1}{2}(S - r) + \frac{\sigma}{2} S\lambda,$$
$$I_{0,\sigma} = \frac{1 - r}{2} - \frac{1}{2}(S - r) + \frac{\sigma}{2} \lambda - \frac{\sigma}{2} S\lambda$$

express each state indicator as its proposed mean plus a linear combination of the three centred tests. Expand the indicator at n first. Every nonconstant term has one of the sources F in (8.1), while the indicator at m is a fixed linear combination of its permitted terminal functions V. Its contribution is therefore $o(1)$, uniformly in A, D. The remaining term is

$$\frac{\pi_i}{2} \, \mathcal{E}_{N,\chi}(A(n) D(m) I_{j,\tau}(m)).$$

Expand the second indicator and apply the reflected version of (8.1), with terminal function 1. This gives (8.2). □

Taking $A = D = 1$ and using the Riemann-sum formula for χ gives

$$\frac{1}{N}\sum_{1\le n<N}\chi(n/N)I_{i,\sigma}(n)I_{j,\tau}(N-n) = \frac{\pi_i\pi_j}{4}\int_0^1 \chi(t)\,dt + o_{\varepsilon,\chi}(1). \tag{8.3}$$

Apply Lemma 4.3 to the bounded product of state indicators, with $c = \pi_i\pi_j/4$. It gives

$$\#\{1 \le n < N : n/N \in I,\ S(n) = i,\ S(N-n) = j, \lambda(n) = \sigma, \lambda(N-n) = \tau\} \\ = |I|\frac{\pi_i\pi_j}{4}N + o_\varepsilon(N), \tag{8.4}$$

uniformly over all intervals $I \subseteq (0,1)$, including $(0,1)$. This proves Theorem 1.2. The uniformity concerns an absolute error divided by N, as specified in the lemma.

Summing over i, j gives the four sign frequencies. Taking $i = j = 1$ gives $|I|r^2N/4$ for each pair of signs among the two smooth summands. Thus the sign and smoothness conclusions are different projections of one finite law; their simultaneous assertion uses the three source tests in (8.1).

8.2 The squarefree local factor

The restriction to nonzero Möbius values is a restriction to squarefree integers. Its two-site density depends on N, since the forbidden residue classes at a prime square may coincide.

The global unsigned asymptotic, including the local factor below, belongs to the classical theory of representations by squarefree integers; see [15, 16] and [17, pp. 591–592]. We include the elementary argument to obtain the weighted form used here. The sign equidistribution in Proposition 8.3 is the additional conclusion.

Lemma 8.2 (Two squarefree summands)**.** *Define*

$$\mathfrak{S}_2(N) = \prod_{p^2\mid N}(1-p^{-2})\prod_{p^2\nmid N}(1-2p^{-2}). \tag{8.5}$$

For every fixed $\chi \in C_c^\infty((0,1))$,

$$\frac{1}{N}\sum_{1\le n<N}\chi(n/N)\mu(n)^2\mu(N-n)^2 = \mathfrak{S}_2(N)\int_0^1\chi(t)\,dt + o_\chi(1). \tag{8.6}$$

The error tends to zero along all integers N.

Proof. First fix a real cutoff $z \ge 2$ and test only squares of primes $p \le z$. Modulo p^2, one excludes $n \equiv 0$ and $n \equiv N$. These are the same residue precisely when $p^2 \mid N$. The Chinese remainder theorem therefore gives the density

$$\mathfrak{S}_{2,z}(N) = \prod_{\substack{p\le z\\ p^2\mid N}}(1-p^{-2})\prod_{\substack{p\le z\\ p^2\nmid N}}(1-2p^{-2})$$

for a periodic mask of period $M_z = \prod_{p\le z}p^2$. Summing the smooth weight in each residue class gives

$$\frac{1}{N}\sum_{1\le n<N}\chi(n/N)\prod_{p\le z}\mathbf{1}_{p^2\nmid n}\mathbf{1}_{p^2\nmid N-n} = \mathfrak{S}_{2,z}(N)\int_0^1\chi(t)\,dt + O_{\chi,z}(N^{-1}).$$

The error is uniform in the residues of N. The population removed by an additional square p^2 at either site is at most $2\lfloor (N-1)/p^2 \rfloor$. The normalized contribution of all primes $p > z$ is thus at most

$$2\|\chi\|_\infty \sum_{p>z} p^{-2}.$$

The same convergent tail controls $|\mathfrak{S}_{2,z}(N) - \mathfrak{S}_2(N)|$, uniformly in N. Let $N \to \infty$ with z fixed, and then let $z \to \infty$. This proves (8.6) without applying the Chinese remainder theorem at the growing cutoff B. □

Proposition 8.3 (Prescribed Möbius signs)**.** *For intervals $I \subseteq (0,1)$, including $I = (0,1)$, and $\sigma, \tau \in \{-1, 1\}$,*

$$\#\{1 \le n < N : n/N \in I,\ \mu(n) = \sigma, \mu(N-n) = \tau\} = \frac{|I|}{4}\mathfrak{S}_2(N)N + o(N), \quad (8.7)$$

$$\#\{1 \le n < N : n/N \in I,\ S(n) = S(N-n) = 1, \mu(n) = \sigma, \mu(N-n) = \tau\} = \frac{|I|r^2}{4}\mathfrak{S}_2(N)N + o_\varepsilon(N). \quad (8.8)$$

Both errors are uniform over all intervals I and sign pairs.

Proof. Put $Q_B(v) = \prod_{p \le B} \mathbf{1}_{p^2 \nmid v}$. This is a function of $v_{\le B}$ with values in $\{0,1\}$, and

$$\frac{1}{N}\#\{1 \le v < N : Q_B(v) \ne \mu(v)^2\} \le \sum_{p>B} p^{-2} = O(B^{-1}).$$

Hence the replacement of the two squarefree masks by $Q_B(n)Q_B(m)$ costs $o(1)$ in normalized counting. On squarefree integers $\mu(v) = \lambda(v)$. Apply Lemma 8.1 with both local masks equal to Q_B and with $i = j = 1$. Then replace Q_B back by μ^2 and apply Lemma 8.2. This proves the smooth formula with the weight χ. Sum Lemma 8.1 over i, j for the formula without smoothness. Apply Lemma 4.3 to both weighted formulas. Their main coefficients depend on N through $\mathfrak{S}_2(N)$, but are bounded by one; the lemma therefore gives both assertions uniformly over intervals. □

Summing (8.7) over the four sign pairs with coefficient $\sigma\tau$ cancels their main terms. Terms with a zero Möbius value contribute zero as well. Consequently

$$\sum_{\substack{1 \le n < N \\ n/N \in I}} \mu(n)\mu(N-n) = o(N), \quad (8.9)$$

uniformly over intervals I. With $I = (0,1)$ this is Conjecture 4.3 of [14]; the almost-all version is recorded in Proposition 4.4 of that paper.

The product (8.5) is a squarefree density. In particular, it is unrelated to the singular series for two primes. Its factors stay uniformly away from a zero product, since

$$\mathfrak{S}_2(N) \ge \prod_p (1 - 2p^{-2}) > 0.$$

The formula is an asymptotic with an N-dependent main term; it does not assert that $\mathfrak{S}_2(N)$ itself converges.

8.3 Functions with a fixed prime budget

The next extension uses ordinary multiplicativity only. Values at prime powers need not be powers of the corresponding prime value. A short-prime factor records those values exactly, while a large prime square is rare.

Lemma 8.4 (Reduction to Liouville with a local factor)**.** *Let $\beta < 1$ and $K \geq 0$ be fixed. Suppose that $f : \mathbb{N} \to \{-1, 1\}$ is multiplicative and*

$$\sum_{p : f(p) \neq -1} p^{-\beta} \leq K. \tag{8.10}$$

For any $B > 1$, put $A_f(u) = f(u)\lambda(u)$. Outside at most

$$N\left(KB^{\beta-1} + \frac{1}{B-1}\right) \tag{8.11}$$

integers $1 \leq n < N$, one has

$$f(n) = \lambda(n)A_f(n_{\leq B}). \tag{8.12}$$

This estimate is uniform in f.

Proof. Discard the integers divisible by an exceptional prime $p > B$ for which $f(p) \neq -1$, and those divisible by p^2 for any prime $p > B$. Their number divided by N is at most

$$\sum_{\substack{p > B \\ f(p) \neq -1}} \frac{1}{p} + \sum_{p > B} \frac{1}{p^2} \leq B^{\beta-1} \sum_{p : f(p) \neq -1} p^{-\beta} + \frac{1}{B-1}.$$

On every remaining integer, the prime factors of $n/n_{\leq B}$ are distinct, exceed B, and have $f(p) = -1$. Multiplicativity gives

$$f(n) = f(n_{\leq B})\lambda(n/n_{\leq B}) = \lambda(n)f(n_{\leq B})\lambda(n_{\leq B}),$$

which is (8.12). No condition on $f(p^k)$ for $k \geq 2$ has been used. □

Proposition 8.5 (Uniform signs under prime perturbations)**.** *Fix $\beta < 1$, $K \geq 0$ and $0 < \varepsilon < 1/4$, and let $I \subseteq (0, 1)$ be an interval, with $I = (0, 1)$ permitted. Let $f, g : \mathbb{N} \to \{-1, 1\}$ be multiplicative functions, each satisfying (8.10). They may differ and may vary with N. Uniformly in these choices and in I, for $i, j \in \{0, 1\}$ and $\sigma, \tau \in \{-1, 1\}$,*

$$\begin{aligned}\#\{1 \leq n < N : n/N \in I,\ S(n) = i,\ S(N-n) = j, f(n) = \sigma,\ g(N-n) = \tau\} \\ = \frac{|I|\pi_i\pi_j}{4}N + o_{\beta,K,\varepsilon}(N).\end{aligned} \tag{8.13}$$

In particular, each pair of signs has $|I|N/4 + o_{\beta,K}(N)$ occurrences when no smoothness condition is imposed.

Proof. Use Lemma 8.4 with $B = \exp(L^{0.26})$ for f at n and for g at $m = N - n$. Reflection is a bijection of $\{1, \dots, N-1\}$, so the exceptional proportion is $O(KB^{\beta-1} + B^{-1}) = o(1)$ uniformly in f, g. Away from it, set $A(n) = A_f(n_{\leq B})$ and $D(m) = A_g(m_{\leq B})$. The sign indicators are

$$\mathbf{1}_{f(n)=\sigma} = \frac{1 + \sigma\lambda(n)A(n)}{2}, \qquad \mathbf{1}_{g(m)=\tau} = \frac{1 + \tau\lambda(m)D(m)}{2}.$$

Multiply by the prescribed smoothness indicators and expand. Every term containing a Liouville sign at n uses source λ or $S\lambda$ in (8.1), with modulation A and terminal function

$1, S, \lambda$ or $S\lambda$, possibly multiplied by D. Such terms are $o(1)$. Terms with a sign only at m are treated after reflection. The unsigned term equals $\pi_i\pi_j \int \chi + o(1)$ by Lemma 8.1 summed over its signs. All coefficients of these finite expansions are bounded independently of f, g, N. This proves the weighted formula uniformly. Lemma 4.3, with the pair (f, g) included in the family index, proves (8.13) uniformly in the interval and the two functions. Summing over i, j proves the final statement; for this sign-only conclusion one may fix, for example, $\varepsilon = 1/8$. □

The fixed condition $\beta < 1$ is used in (8.11). This argument supplies no uniform extension at the reciprocal endpoint $\beta = 1$. The case $f = g = \lambda$ is included, as is the completely multiplicative class in the earlier all-total lower-bound theorem. Thus the uniform conclusion here strengthens that earlier counting conclusion as well as its Liouville special case.

A concrete example of ordinary multiplicativity is the unitary Möbius function

$$g(n) = (-1)^{\omega(n)},$$

where $\omega(n)$ counts distinct prime factors. It satisfies $g(p) = -1$ at every prime, so its prime budget is zero. However, $g(p^2) = -1 \neq g(p)^2$, and it is not completely multiplicative. Taking $f = g$ in Proposition 8.5 gives each of its reflected sign pairs frequency $1/4 + o(1)$ at every total, together with the joint smoothness law; mixed pairs with $f = \lambda$ are also allowed. Tian and Yang prove lower bounds for the same unitary sign pairs at every sufficiently large total and quarter frequencies at almost all totals [14, Theorems 1.5 and 1.7]. Their exact classification at small totals is given separately in Theorem 1.6.

8.4 Positive coefficients in a linear equation

A divisibility constraint at each end is a local test. For a uniform *absolute* estimate all large coefficients can also be treated directly by counting solutions.

Corollary 8.6 (Positive linear equations)**.** *Assume the F_λ instance of (2.11). Then*

$$\sup_{\substack{a,b\geq 1\\ I\subseteq(0,1)}} \frac{1}{N}\left|\sum_{\substack{j,k\geq 1\\ aj+bk=N\\ aj/N\in I}} \lambda(j)\lambda(k)\right| \longrightarrow 0, \tag{8.14}$$

where a, b are integers and I is an interval. For each fixed smooth interior weight $\chi(aj/N)$ in place of the interval indicator, the normalized sum is $O_\chi((\log N)^{-1/1000})$, uniformly in a, b.

Proof. Fix $\varepsilon = 1/8$ in the sign-only construction. If $\max(a, b) \leq B$, both $\mathbf{1}_{a|u}$ and $\mathbf{1}_{b|v}$ are functions of $u_{\leq B}$ and $v_{\leq B}$ respectively. Put $u = aj$, $v = bk$. Complete multiplicativity introduces the constant $\lambda(a)\lambda(b)$. Proposition 4.2 bounds the resulting normalized smooth sum by $O_\chi(L^{-1/1000})$ uniformly in the two masks. If $\max(a, b) > B$, there are at most $N/\max(a, b)$ solutions: fixing the variable with the larger coefficient determines the other one. The normalized smooth sum is then at most $\|\chi\|_\infty/B$. This proves the weighted assertion. For the interval claim apply Lemma 4.3 with (a, b) as the family index and constant zero, extending each summand by zero off its divisibility masks. □

The denominator in (8.14) is N. For $N > 2$, the choice $a = N - 1$, $b = 1$ has the single positive solution $(1, 1)$, so its mean sign product per representation is one.

8.5 Different smoothness thresholds at the two summands

Two cutoffs ask a different question even when the signs are omitted. For example, imposing $P^+(n) \le N^{1/12}$ and $P^+(N-n) \le N^{1/6}$ describes a rectangle in the two largest prime factors. Knowing only the mass of all squares with a common cutoff would leave the mass of this rectangle undetermined.

Corollary 8.7 (Two smoothness thresholds)**.** *Assume the marked estimate for the three sources at each fixed exponent in $(0,1/4)$. For $0 \le t \le 1/4$, put $S_t(v) = \mathbf{1}_{P^+(v)\le N^t}$, $r(t) = \rho(1/t)$ for $t > 0$, and $r(0) = 0$. Write $\pi_t(1) = r(t)$ and $\pi_t(0) = 1 - r(t)$. For $i, j \in \{0,1\}$ and $\sigma, \tau \in \{-1,1\}$ one has*

$$\#\{1 \le n < N : n/N \in I,\ S_s(n) = i,\ S_t(N-n) = j,\ \lambda(n) = \sigma,\ \lambda(N-n) = \tau\} = \frac{N|I|}{4}\,\pi_s(i)\pi_t(j) + o(N).$$

The error divided by N tends to zero uniformly over both thresholds $s, t \in [0,1/4]$, all intervals $I \subseteq (0,1)$ and all sixteen states.

Proof. Fix $0 < s, t < 1/4$ first. The prime groups and the cutoffs B, Q in (2.1) depend only on L and T. For large N, $Q < N^{\min(s,t)}$. Use the two centres of Section 5 with $y = N^s$ at the source. Their covariance and Proposition 5.3 give the inputs of Proposition 7.1 at that site. The target reader determines both $S_t(v)$ and $\lambda(v)$ from v° and the group counts. Indeed, every omitted prime is smaller than N^t. The same statement holds for the coarser target observation of Proposition 3.4. Its inversion cost is independent of s, t. On writing $n = du$, $N - n = dv$, the retained $d \le G$ has only prime factors below both cutoffs. Thus the proof of Proposition 4.2 applies to sources $\lambda, S_s - r(s), S_s\lambda$ against $1, S_t, \lambda, S_t\lambda$. To interchange the sites use the centres for $y = N^t$ and reflect the test χ. Take L above both lower bounds required by these applications, with s and t fixed. The indicator expansion in Lemma 8.1 gives $\pi_s(i)\pi_t(j)/4$; Lemma 4.3 gives the fixed-(s,t) conclusion uniformly in I.

We next pass to variable thresholds without taking a supremum in a graph or centre estimate. At one site let

$$D_N(x) = N^{-1}\#\{1 \le n < N : P^+(n) \le N^x\}.$$

Proposition 5.4, summed over the two signs with multiplier one, gives $D_N(x) \to r(x)$ throughout $[0,1/4]$, including the endpoints. Reflection permutes $\{1, \dots, N-1\}$, so the same marginal bound applies at the second site.

Choose a finite partition of $[0,b]$ with at least one interior point and with r-increment at most η in each cell. Each threshold can be replaced by an interior endpoint of its cell. The symmetric difference between the two joint events before and after replacement is contained in the union of the two marginal bands, for each choice of i, j, σ, τ and I. Each band has normalized size at most $\eta + o(1)$ by the marginal convergence at the finitely many endpoints. The limiting expression changes by at most 2η. The fixed-interior formula applies simultaneously at the finitely many chosen pairs. Taking $N \to \infty$ and then $\eta \downarrow 0$ proves the uniform assertion. □

For the rectangle in the opening example, the proportion for each prescribed pair of signs on the full interval tends to $\rho(12)\rho(6)/4$. If a threshold tends to zero, the corollary still controls the discrepancy divided by N. When the predicted proportion tends to zero, an asymptotic relative to that proportion would require a different estimate.

9 Range of the conclusions

The one-site result, Proposition 5.4, uses the classical prime number theorem. In the two-site problem the parameter that tends to infinity is the total N. The reflection point therefore moves with the sum. All counts in this paper are ordinary counts of ordered representations, and the reflected conclusions of Theorems 1.1 and 1.2 are conditional on the marked estimate along the full sequence of positive integers. The common-divisor argument is what permits that last assertion: a result restricted to prime or rough totals would not suffice for it.

The analytic comparison is made at fixed positive smoothness exponents. Corollary 8.7 passes from this comparison to a uniform absolute error for two thresholds. The positive constant $\rho(1/\varepsilon)^2/4$ can be very small, but for each fixed ε it gives a positive limiting proportion in every fixed interval of positive length. Neither the centre estimates nor the graph parameter choices here are uniform as $\varepsilon \downarrow 0$. Lemma 4.3 does give a uniform absolute $o(N)$ error over moving intervals, but no relative error $o(N|I_N|)$ for shrinking intervals. Corollary 8.6 permits all positive integer coefficients when the error is measured against N. For a single representation its contribution has size one, whatever the value of N. For the multiplicative extension, the fixed gap $1-\beta>0$ is used explicitly in the exceptional-density estimate (8.11).

Under the marked premise the logarithmic error in the smooth local estimate is sufficient to remove the occupancy weights and sum the common-divisor strata. In passing to sharp interval counts, we fix the smooth approximation and a finite grid first and take their limits afterwards. Consequently the main theorems assert a uniform absolute $o(N)$ error without an explicit rate for those sharp counts. We also make no numerical claim about the first total from which an asymptotic lower bound holds.

The squarefree factor retains the residue dependence on N which survives in the Möbius problem. The Liouville signs themselves have no such factor in the limiting law. The local masks in Section 8 separate these two effects: the elementary density of the masks provides the local factor, while the signed estimates provide the factor $1/4$.

The graph comparison uses the zero-free half-plane for all Dirichlet characters stated in [8, Theorem 1.1] and imported in [7, Theorem 3.2]. Its role is the uniform selected-prime estimate at moduli exponentially larger than the selected primes. The one-variable centering argument of Section 5 uses only the classical prime number theorem. These two analytic inputs have different strengths and have been kept separate throughout the proof.

The four forms in the author's *An explicit bound below one half for a quartic Liouville correlation* [23] are $n, n+1, 3n+1, 3n+2$. For the first pair the total is $2n+1$, and for the second it is $6n+3$. Each reflected average here runs over all first summands of a prescribed total. Selecting the single summand n of $2n+1$ would require an estimate on a diagonal. Even uniform absolute $o(N)$ errors on intervals permit arbitrary behaviour at an individual summand. In particular the reflected law supplies no estimate for the simultaneous product of those four signs.

9.1 What a smoothed reflected sum detects

A simple example explains why estimates for averages of convolution coefficients lead to a different problem. Let χ_3 be the nonprincipal character modulo 3 and define the completely multiplicative sign

$$f(n)=\chi_3(n/3^{v_3(n)}), \qquad C_f(N)=\sum_{1\le n<N} f(n)f(N-n).$$

The partial sums of χ_3 have absolute value at most one. Splitting into powers of 3 gives $A_f(x)=\sum_{n\le x} f(n)=O(\log(2x))$ for $x\ge 1$. Integrating the two partial sums and computing

the length contributed by each ordered pair gives, for $x \geq 2$,

$$\sum_{N \leq x} (x - N) C_f(N) = \int_0^x A_f(u) A_f(x - u)\, du = O(x \log^2(2x)),$$

with $A_f(u) = 0$ for $u < 1$. Yet at $N = 3^k$, if $n = 3^j m$ with $3 \nmid m$ and $n < N$, the unit part of $N - n$ is $3^{k-j} - m \equiv -m \bmod 3$. Hence $C_f(3^k) = -(3^k - 1)$ for every $k \geq 1$. The mechanism is the residue of the unit part at a fixed prime. It shows why a Cesàro bound of this quality by itself cannot yield cancellation at each total, even for a completely multiplicative sign. The arithmetic content specific to Liouville enters our argument in the proposed marked comparison.

Assistance and verification

Brice Pouly is the author. The research was developed with substantial assistance from GPT-6 Astra. GPT-6 Astra Pro assisted with mathematical derivations, the assembly of the graph and counting arguments, bibliographic checking and English exposition. An OpenAI model used through Codex assisted with the one-site and two-threshold laws, the colour-averaging calculation, the analysis of the graph interfaces and the finite checks supplied with the Zenodo sources.

The proof of the marked estimate is a candidate argument. The finite checks concern the colour identities, the omission congruence and the character-energy normalization on specified finite instances. The verification archive of [7] concerns the statements in that work. The extension in Sections 6 and 7 has received model-assisted scrutiny. A full independent human or proof-assistant verification of that extension remains open.

References

[1] C. A. Corrádi and I. Kátai, *Some problems concerning the convolutions of number-theoretical functions*, Archiv der Mathematik **20** (1969), 24–29. doi:10.1007/BF01898987.

[2] A. P. Mangerel, *On a Goldbach-type problem for the Liouville function*, International Mathematics Research Notices **2024** (2024), no. 16, 11865–11877. doi:10.1093/imrn/rnae149.

[3] A. P. Mangerel, *On Shusterman's Goldbach-type problem for sign patterns of the Liouville function*, preprint, arXiv:2412.17199v1, 23 December 2024. doi:10.48550/arXiv.2412.17199.

[4] L. Meng, *All four Liouville sign patterns occur in the binary decompositions of every integer $N \geq 11$*, preprint, eScholarship, 2026, item 7945n3rh. Author's manuscript.

[5] K. Krishnamoorthy, *On a conjecture of Corrádi and Kátai*, preprint, arXiv:2608.13266v1, 13 August 2026. doi:10.48550/arXiv.2608.13266.

[6] B. Pouly, *Reflected sign patterns of the Liouville function at every large total*, Zenodo, 2026. doi:10.5281/zenodo.22959603. Also available as hexagon:2610.00350v1.

[7] W. Howe, G. Huffman, J. Kohlmeyer and L. Kruer, *Two smooth summands for every large integer*, preprint, manuscript dated 13 November 2026, hexagon:2611.00787v1, 21 November 2026. Pinned source and verification archive, commit `0f66e1d94275854de3a6bee7d109c4d2b5cecb35`.

[8] OpenAI, *The Quasi-Riemann Hypothesis. A Zero-Free Half-Plane $\Re s > 7/8$*, preprint, 30 September 2026. Pinned manuscript, commit `adc7f1241b42e322a6451854ab7e4b4c146bf78a`.

[9] OpenAI, *Weighted dilation graphs, smooth shifted primes and totient fibers*, preprint, 24 September 2026. Pinned manuscript, commit `4eaee87548250b2682affda7d05d947c05554bae`.

[10] H. Davenport, *Multiplicative Number Theory*, third edition, revised by H. L. Montgomery, Graduate Texts in Mathematics, vol. 74, Springer, New York, 2000. ISBN 978-0-387-95097-6.

[11] A. Hildebrand and G. Tenenbaum, *Integers without large prime factors*, Journal de théorie des nombres de Bordeaux **5** (1993), no. 2, 411–484. doi:10.5802/jtnb.101. Digitized article.

[12] J.-M. De Koninck, L. Germán and I. Kátai, *On the convolution of the Liouville function under the existence of Siegel zeros*, Lithuanian Mathematical Journal **55** (2015), no. 3, 331–342. doi:10.1007/s10986-015-9284-x.

[13] M. Cantarini, A. Gambini and A. Zaccagnini, *On the discrete convolution of the Liouville and Möbius functions*, preprint, arXiv:2603.10241v1, 10 March 2026. doi:10.48550/arXiv.2603.10241.

[14] Tian Caolin and Yang Chenxi, *A Goldbach-Type Problem for the Unitary Möbius Function*, manuscript dated 2 November 2026, arXiv:2611.10499v1, 7 November 2026. doi:10.48550/arXiv.2611.10499.

[15] E. H. Linfoot and C. J. A. Evelyn, *On a problem in the additive theory of numbers*, Mathematische Zeitschrift **30** (1929), 433–448. doi:10.1007/BF01187781.

[16] T. Estermann, *On the representations of a number as the sum of two numbers not divisible by k-th powers*, Journal of the London Mathematical Society **s1-6** (1931), no. 1, 37–40. doi:10.1112/jlms/s1-6.1.37.

[17] J. Brüdern and A. Perelli, *Exponential sums and additive problems involving square-free numbers*, Annali della Scuola Normale Superiore di Pisa, Classe di Scienze, series 4, **28** (1999), no. 4, 591–613. Digitized article.

[18] J. Teräväinen, *On binary correlations of multiplicative functions*, Forum of Mathematics, Sigma **6** (2018), e10. doi:10.1017/fms.2018.10. The version cited for Theorem 1.14 is arXiv:1710.01195v2, 21 May 2018.

[19] OpenAI, *The joint Dickman law for consecutive integers*, preprint, 24 September 2026. Pinned manuscript, commit `4eaee87548250b2682affda7d05d947c05554bae`.

[20] OpenAI, *Ordinary two-point correlations of multiplicative functions*, preprint, 24 September 2026. Pinned manuscript, commit `4eaee87548250b2682affda7d05d947c05554bae`.

[21] A. Balog, *On additive representation of integers*, Acta Mathematica Hungarica **54** (1989), 297–301. doi:10.1007/BF01952060.

[22] S. Jain, *Sums of Integers with Small Prime Factors*, preprint, arXiv:2611.21136v1, 16 November 2026. doi:10.48550/arXiv.2611.21136.

[23] B. Pouly, *An explicit bound below one half for a quartic Liouville correlation*, Zenodo, 2026. doi:10.5281/zenodo.24443184. hexagon:2611.00891v1.

[24] L. P. Slothouber, *Phase cones for Gamma-weighted atoms of additive Dirichlet series*, 26 November 2026. Hexagon 2611.01111v1.

[25] B. Pouly, *Joint cofactor signs in prime-factor windows on consecutive integers*, Zenodo, 2026. doi:10.5281/zenodo.24578855.

[26] A. Hildebrand, *On the number of prime factors of integers without large prime divisors*, Journal of Number Theory **25** (1987), 81–86. doi:10.1016/0022-314X(87)90017-5.

[27] R. de la Bretèche and G. Tenenbaum, *Friable averages of oscillating arithmetic functions*, in H. Maier, J. Steuding and R. Steuding (eds.), *Number Theory in Memory of Eduard Wirsing*, Springer, 2023, 43–71. doi:10.1007/978-3-031-31617-3_5.