A Proposed Solution to Erdős Problem 390

Shouqiao Wang

Columbia University Multiscalar Intelligence

Abstract

Let $f(n)$ be the least possible largest factor in a representation of $n!$ as a product of distinct integers, all larger than n. We propose the exact asymptotic

$$\lim_{n\to\infty} \frac{(f(n)-2n)\log n}{n} = \frac{4029639598}{25970038185}.$$

The lower bound is a thirteen-layer valuation cut. The upper bound uses an exact cofactor-allocation certificate, central-binomial anchors, a guarded rough-signature selector, a marked friable-number covariance bridge, a finite-band tangent correction, and deterministic column-sparse rounding with a precharged universal switch bank. The complete finite allocation certificate and its exact checker are supplied in the companion file `numerical_verifier.py`. This proposed solution was found by GPT-5.6.

Contents

1 Introduction

For an integer $n \geq 3$, let $f(n)$ be the least possible value of the largest factor in a factorization

$$n! = a_1 a_2 \cdots a_k, \qquad n < a_1 < a_2 < \cdots < a_k.$$

Thus the factors are required both to be distinct and to lie strictly above n. Erdős, Guy and Selfridge proved that there are absolute constants $0 < c_1 < c_2$ for which

$$2n + c_1 \frac{n}{\log n} < f(n) < 2n + c_2 \frac{n}{\log n}$$

for all sufficiently large n, and they asked whether the second-order term has an asymptotic constant [EGS82, Theorem 3].

Subsequent discussion isolated two features that are useful here. Tao emphasized the complementary product formulation below and observed that the original lower-bound count misses further carry intervals [Tao25]. Mausberg made the first thirteen of these intervals quantitative and obtained the lower-bound constant displayed in (1.1) [Mau26]. The lower-bound argument in this paper is based on that thirteen-layer obstruction, with an additional reduction that treats endpoints at or below $2n$. We make no claim here about priority beyond these cited records. For a recent treatment of the related, but different, problem of decomposing a factorial into a prescribed number of large factors, see Alexeev et al. [ACR+26]; none of its quantitative estimates is used below.

Define

$$C_0 := \frac{\displaystyle\sum_{r=1}^{13} \frac{1}{(r+1)(2r+1)}}{\displaystyle\sum_{\substack{p \le 23 \\ p \text{ prime}}} \frac{1}{p-1}} = \frac{4029639598}{25970038185} = 0.15516494697830188\ldots. \tag{1.1}$$

Our main result determines the conjectured constant.

Theorem 1.1. *As $n \to \infty$,*

$$f(n) = 2n + C_0 \frac{n}{\log n} + o\left(\frac{n}{\log n}\right).$$

Equivalently,

$$\lim_{n \to \infty} \frac{(f(n) - 2n)\log n}{n} = C_0.$$

The proof is most naturally expressed after taking complements. For an integer $M > n$, put

$$Q(n, M) := \frac{M!}{(n!)^2}, \qquad \mathcal{I}(n, M) := (n, M] \cap \mathbb{Z}.$$

Lemma 1.2 (Complement formulation)**.** *There is a factorization of $n!$ into distinct members of $\mathcal{I}(n, M)$ if and only if $Q(n, M)$ is the product of a subset of the distinct integers in $\mathcal{I}(n, M)$.*

Proof. The product of every integer in $\mathcal{I}(n, M)$ is $M!/n!$. If one subset has product $n!$, its complement therefore has product $(M!/n!)/n! = Q(n, M)$. The same calculation in reverse proves the converse. □

We briefly describe the proof of Theorem 1.1. The lower bound uses thirteen disjoint layers of primes P for which $v_P(Q(n, M)) = 1$. The unique selected multiple of such a prime has the form Pq, with $r + 1 \le q \le 2r + 1$ in layer r. Every cofactor arising in the first thirteen layers has a prime divisor at most 23. Comparing these forced incidences with the available small-prime valuations gives exactly the quotient in (1.1).

For the upper bound, fix $c > C_0$, set

$$L = \log n, \qquad N = \frac{n}{L}, \qquad M = 2n + \lceil cN \rceil,$$

and construct a distinct subset of $(n, M]$ with product $Q(n, M)$. The construction has five interacting parts. First, an exact rational allocation of the large-prime carry layers produces central anchors; its finite portion is certified by rational arithmetic and its infinite tail by Nagura's prime-interval theorem. Second, all factors needed by the integral repair are reserved in advance, so that their valuation cost is present in the target ledger before any equation is solved. Third, with $y = n^{2/9}$, a signed selector matches every complete y-rough signature exactly and leaves a residual $O(N/(pL))$ at each prime $p \in (W, y]$. Fourth, marked smooth-number estimates and the Poisson–Dickman covariance operator fit the smooth row, the fixed head valuations, the physical logarithm and finitely many exponent bands. Finally, local same-row switches cancel the remaining medium-prime residual, and a precharged universal bank converts the fractional point to a 0-1 point without changing any exact row quota. The resulting exact valuation identity is the product identity required by Lemma 1.2.

The role of the computer-assisted component is deliberately narrow. It checks a finite list of rational equalities and inequalities for the anchor allocation. The complete finite certificate and its exact checker are the single companion file `numerical_verifier.py`; Appendix A records its precise scope. The analytic tail argument is proved in Section 4.1, and no floating-point output is used as a premise of the proof.

2 Analytic and probabilistic inputs

This section records the external results used later, together with the specializations needed in the proof. All limiting parameters appearing below are fixed before $n \to \infty$. In particular, constants may depend on fixed compact intervals, the fixed head cutoff and a fixed finite cell partition, but never on n or on a marked prime.

2.1 Smooth numbers and the de Bruijn–Saias normalization

Write $P^+(m)$ and $P^-(m)$ for the largest and smallest prime factors of m, with $P^+(1) = 1$, and put

$$\Psi(x, y) := \#\{m \le x : P^+(m) \le y\}, \qquad u := \frac{\log x}{\log y}.$$

The Dickman function is defined by

$$\begin{aligned} \rho(u) &= 0 \quad (u < 0), \\ \rho(u) &= 1 \quad (0 \le u \le 1), \\ u\rho'(u) + \rho(u-1) &= 0 \quad (u > 1). \end{aligned}$$

At transition points it is important not to replace the smooth-number count prematurely by $x\rho(u)$. We use the de Bruijn–Saias normalization

$$\Lambda(x, y) := x \int_{-\infty}^{\infty} \rho(u - v)\, d(\lfloor y^v \rfloor y^{-v})\,, \tag{2.1}$$

where at integral x the value means the right limit $\Lambda(x + 0, y)$. This one-sided convention is part of the definition in the cited source.

Theorem 2.1 (Hildebrand–Tenenbaum–Saias)**.** *For every $\varepsilon > 0$, uniformly in*

$$y \ge y_0(\varepsilon), \qquad 1 \le u \le \exp\{(\log y)^{3/5-\varepsilon}\},$$

one has

$$\Psi(x, y) = \Lambda(x, y) \left\{1 + O_\varepsilon\!\left(e^{-(\log y)^{3/5-\varepsilon}}\right)\right\}. \tag{2.2}$$

On every fixed compact u-range contained in $[1, \infty)$,

$$\Lambda(x, y) = x\rho(u) \left\{1 + O\!\left(\frac{\log(u+1)}{\log y}\right)\right\}. \tag{2.3}$$

For $0 < u \le 1$, the exact identity $\Psi(x, y) = \lfloor x \rfloor$ is used instead.

The formulation is Theorem 1.8, (1.29) and Lemma 3.1 of Hildebrand–Tenenbaum [HT93]; the underlying sharp estimate is due to Saias [Sai89]. In our application $y = n^{2/9} \to \infty$, while all values of u remain in a fixed compact set. Thus both range conditions in Theorem 2.1 hold with room to spare, and the exponential error is smaller than every fixed negative power of $\log n$. Formula (2.2), rather than only (2.3), will be retained at terminal and Dickman transition points.

We shall repeatedly use the following derived marked form. It is included here to specify exactly what is, and what is not, obtained from the cited smooth-number theorem.

Lemma 2.2 (Uniform marked smooth cells)**.** *Let $y = n^{2/9}$, so that $U = \log n/\log y = 9/2$, and put $L = \log n$. Fix a finite set H of primes, a vector $\mathbf{e} = (e_\ell)_{\ell\in H}$ of nonnegative integers, and fixed reals $0 < A < B < \infty$. Define*

$$h_{\mathbf{e}} := \prod_{\ell\in H} \ell^{e_\ell}, \qquad M_H := \prod_{\ell\in H} \ell, \qquad \delta_{\mathbf{e}} := \frac{1}{h_{\mathbf{e}}}\frac{\varphi(M_H)}{M_H}.$$

If $d \le y^4$ is y-smooth and $(d, M_H) = 1$, then, uniformly in d,

$$\begin{aligned}\#\{m : An < m \le Bn,\ P^+(m) \le y,\ d \mid m,\ v_\ell(m) = e_\ell\ (\ell \in H)\}\\ = \delta_{\mathbf{e}}\frac{(B-A)n}{d}\left\{\rho\Big(U - \frac{\log d}{\log y}\Big) + O_{H,A,B}\Big(\frac{1}{L}\Big)\right\}.\end{aligned} \tag{2.4}$$

If ϕ has bounded variation on $[A, B]$, the corresponding Stieltjes-weighted form is

$$\begin{aligned}\sum_{\substack{An<m\le Bn,\ P^+(m)\le y,\ d\mid m\\ v_\ell(m)=e_\ell\ (\ell\in H)}} \phi(m/n)\\ = \delta_{\mathbf{e}}\frac{n}{d}\rho\Big(U - \frac{\log d}{\log y}\Big)\int_A^B \phi(t)\,dt + O_{H,A,B}\Big(\frac{n}{dL}\{\|\phi\|_\infty + \operatorname{Var}\phi\}\Big).\end{aligned} \tag{2.5}$$

Consequently, under normalized counting measure on the unmarked cell,

$$\mathbb{P}(d \mid m) = \frac{1}{d}\frac{\rho(U - \log d/\log y)}{\rho(U)} + O_{H,A,B}\Big(\frac{1}{dL}\Big). \tag{2.6}$$

The conclusions persist under a fixed finite disjoint union, or a compactly weighted fixed finite mixture, of such cells. An overlapping finite family is always first disjointified and then interpreted through the induced convex mixture.

Proof. The exact valuation conditions are finite divisibility inclusion–exclusion:

$$1_{\{v_\ell(m)=e_\ell\ (\ell\in H)\}} = 1_{h_{\mathbf{e}}\mid m}\sum_{a\mid M_H}\mu(a)1_{a\mid m/h_{\mathbf{e}}}.$$

After also dividing by d, each summand is a difference of two values of $\Psi(\,\cdot\,, y)$. The fixed factors $A, B, h_{\mathbf{e}}, a$ alter the smoothness parameter by $O_H(1/L)$, and ρ is Lipschitz on the resulting compact range. Theorem 2.1, with the exact formula below parameter one, therefore gives (2.4) after summing

$$\frac{1}{h_{\mathbf{e}}}\sum_{a\mid M_H}\frac{\mu(a)}{a} = \delta_{\mathbf{e}}.$$

The endpoint floor error is $O(1)$. Since $n/d \ge n/y^4 = n^{1/9}$, it is absorbed by $O(n/(dL))$. Stieltjes partial summation proves (2.5); division by the case $d = 1$ proves (2.6). □

Only exact valuation and coprimality patterns, and finite disjoint unions of them (after disjointification when necessary), will be called *head cells*. Lemma 2.2 makes no assertion about arbitrary nonzero residue classes modulo a fixed integer; such a statement would require a smooth-number theorem in arithmetic progressions. The bound $d \le y^4$ covers the marks p, pq, p^k, and p^kq used below and leaves the genuine cofactor margin $n/y^4 = n^{1/9}$.

2.2 Prime counts and Mertens estimates

We use the classical zero-free-region prime number theorem in the form

$$\pi(x) = \mathrm{li}(x) + O\Big(xe^{-c\sqrt{\log x}}\Big) \qquad (x \geq 2) \tag{2.7}$$

for an absolute $c > 0$; see [MV07, Theorem 6.9]. In particular, if $0 < a < b$ are fixed and $n^a \leq X \leq Z \leq n^b$, subtraction at the two exact endpoints gives

$$\pi(Z) - \pi(X) = \int_X^Z \frac{dt}{\log t} + O_{a,b}\Big(Ze^{-c_{a,b}\sqrt{\log n}}\Big). \tag{2.8}$$

The error in (2.8) is absolute. We do not invoke a relative prime number theorem for arbitrary short intervals. Whenever a relative conclusion is used later, the displayed main term is separately checked to dominate this absolute error.

Mertens' estimates, in the forms used here, are

$$\sum_{p \leq x} \frac{1}{p} = \log\log x + B_1 + O\Big(\frac{1}{\log x}\Big), \tag{2.9}$$

$$\prod_{p \leq x} \Big(1 - \frac{1}{p}\Big)^{-1} = e^{\gamma} \log x + O(1).$$

See [Mer74] or [MV07, Theorem 2.7(d),(e)]. Thus, for fixed $0 < a < b$,

$$\sum_{y^a < p \leq y^b} \frac{1}{p} = \log \frac{b}{a} + O_{a,b}\Big(\frac{1}{\log y}\Big), \tag{2.10}$$

and

$$\prod_{p \leq y} \Big(1 - \frac{1}{p}\Big) = \frac{e^{-\gamma}}{\log y} + O\Big(\frac{1}{\log^2 y}\Big).$$

These formulas justify every bounded reciprocal-prime mass used in a fixed-depth expansion. They also give, for $X \geq y$,

$$\sum_{\substack{R \leq X \\ P^-(R) > y}} \frac{1}{R} \leq \prod_{y < p \leq X} \Big(1 - \frac{1}{p}\Big)^{-1} \ll \frac{\log X}{\log y}.$$

For $y = n^{2/9}$ and $X \leq n^C$, this is $O_C(1)$.

2.3 The interval Selberg sieve

For a positive integer P, define

$$S(x, H; P) := \#\{m : x < m \leq x + H,\ (m, P) = 1\}, \qquad L_P(R) := \sum_{\substack{d \leq R \\ d \mid P}} \frac{\mu(d)^2}{\varphi(d)}.$$

The one-dimensional Selberg upper-bound sieve gives, uniformly in the real translate x,

$$S(x, H; P) \leq \frac{H}{L_P(R)} + O\bigg(\frac{R^2}{L_P(R)^2}\bigg) \qquad (H > 0,\ R \geq 1).$$

This is [MV07, Theorem 3.2]. Indeed, for weights supported on $d, e \leq R$, the elementary interval remainder is

$$\#\{m : x < m \leq x + H,\ [d, e] \mid m\} = \frac{H}{[d, e]} + O(1).$$

Take $P = \prod_{p \le y} p$ and choose Selberg weights supported on $d \le R = y^2$. The elementary lower bound

$$L_P(y^2) \ge \sum_{d \le y} \frac{\mu(d)^2}{\varphi(d)} \ge \sum_{d \le y} \frac{\mu(d)^2}{d} = \frac{1}{\zeta(2)} \log y + O(1) \gg \log y$$

follows, for example, by inserting $\mu(d)^2 = \sum_{a^2 | d} \mu(a)$ in the last sum. Consequently,

$$\#\{m : x < m \le x + H : P^-(m) > y\} \ll \frac{H}{\log y} + \frac{y^4}{(\log y)^2}. \tag{2.11}$$

The weight support is y^2, whereas the largest lcm remainder modulus is $[d, e] \le y^4$; we keep these two levels distinct. In the exceptional row application, $y = n^{2/9}$ and $H \gg n^{1-\delta_*}/\log n$, where $\delta_* < 1/18$. Hence $y^4 = n^{8/9} = o(H)$, even after logarithmic factors, and (2.11) is $O(H/\log y)$ uniformly in the translated interval.

2.4 The Poisson–Dickman bridge

Let Π be a Poisson point process on $(0, 1]$ with intensity dt/t, and set

$$T := \sum_{t \in \Pi} t.$$

Arratia, Barbour and Tavaré proved that T has density

$$g(t) = e^{-\gamma} \rho(t) \qquad (t > 0) \tag{2.12}$$

[ABT99, (2.3)–(2.6)]. In particular, $\int_0^\infty \rho(t)\, dt = e^\gamma$.

We fix a canonical conditional version at every exact total. For $0 < \varepsilon < 1$, write

$$\Pi_{>\varepsilon} := \Pi \cap (\varepsilon, 1], \qquad T_{>\varepsilon} := \sum_{x \in \Pi_{>\varepsilon}} x,$$

and let $T_{\le\varepsilon}$ be the independent mass of the remaining atoms. After scaling by ε, the process on $(0, \varepsilon]$ is again the scale-invariant process on $(0, 1]$. Consequently $T_{\le\varepsilon}$ has the continuous density

$$g_\varepsilon(v) = \varepsilon^{-1} g(v/\varepsilon) \qquad (v > 0).$$

If G is a bounded cylinder function depending only on $\Pi_{>\varepsilon}$, define

$$\mathbb{E}_u G := \frac{1}{g(u)} \mathbb{E}\Big[G(\Pi_{>\varepsilon}) g_\varepsilon(u - T_{>\varepsilon}) \mathbf{1}_{\{T_{>\varepsilon} < u\}}\Big]. \tag{2.13}$$

The convolution identity for the density of $T_{>\varepsilon} + T_{\le\varepsilon}$ shows that this is a probability law. Here is the consistency calculation. If $0 < \varepsilon' < \varepsilon$ and G depends only on $\Pi_{>\varepsilon}$, condition first on $\Pi_{>\varepsilon}$ in the ε'-formula. With V the independent mass on $(\varepsilon', \varepsilon]$, the convolution identity is

$$\mathbb{E}\Big[g_{\varepsilon'}(v - V) \mathbf{1}_{\{V < v\}}\Big] = g_\varepsilon(v) \qquad (v > 0).$$

Substitution of $v = u - T_{>\varepsilon}$ shows that the two definitions of $\mathbb{E}_u G$ coincide. The configuration space of locally finite point measures on $(0, 1]$, with the sigma-field generated by the restrictions to $(\varepsilon, 1]$, is standard Borel. The consistent finite restrictions therefore define a unique probability law $\mathbb{P}_u$ for every $u > 0$. For each bounded cylinder function, the right-hand side of (2.13) is a Borel function of u. A monotone class argument, starting from the cylinder events, therefore shows that $u \mapsto \mathbb{P}_u(A)$ is Borel for every configuration event A. Thus $(\mathbb{P}_u)_{u>0}$ is a probability kernel, as required for the disintegrations below.

This projective law is supported on configurations of total mass exactly u, rather than merely at most u. Indeed its restrictions give $T_{>\varepsilon} \le u$, while the finite-ε conditional Campbell calculation gives directly

$$\mathbb{E}_u\{u - T_{>\varepsilon}\} = \int_0^\varepsilon \frac{\rho(u-s)}{\rho(u)}\, ds \le C_u \varepsilon. \tag{2.14}$$

Thus $u - T_{>\varepsilon} \to 0$ in $L^1(\mathbb{P}_u)$, and monotone convergence yields $\sum_{x\in\Pi} x = u$ almost surely. This selects the continuous regular conditional version used below, rather than leaving its value at a probability-zero event arbitrary.

Let $\mathbb{E}_U$ denote expectation under $\mathbb{P}_U$. Campbell–Mecke and its multivariate form [LP17, Theorems 4.1 and 4.4], applied first to cylinder functions in (2.13), and then monotone convergence, gives after division by the density in (2.12), for bounded measurable test functions for which the integrals converge,

$$\mathbb{E}_U \sum_{t\in\Pi} F(t) = \int_0^1 F(s) \frac{\rho(U-s)}{\rho(U)} \frac{ds}{s}, \tag{2.15}$$

$$\mathbb{E}_U \sum_{\substack{s,t\in\Pi \\ s\ne t}} F(s)G(t) = \int_0^1\int_0^1 F(s)G(t) \frac{\rho(U-s-t)}{\rho(U)} \frac{ds}{s}\frac{dt}{t}. \tag{2.16}$$

Here $\rho(v) = 0$ for $v < 0$. Since $g(U) > 0$ for every $U > 0$, these density ratios define the required conditional version at the fixed value $U = 9/2$. In particular, with

$$h(s) := \frac{\rho(U-s)}{\rho(U)}, \qquad K(s,t) := \frac{\rho(U-s-t)}{\rho(U)} - h(s)h(t),$$

equations (2.15)–(2.16) give exactly the multiplication and covariance kernels used by the smooth-row bridge. The formulas hold first for nonnegative tests; a signed test follows by positive and negative parts whenever the right side is absolutely integrable. In every later use $F(t) = O(t)$, or a product of two such functions, so (2.14) and $|K(s,t)| \ll st$ verify the required first and second integrability.

2.5 An asymmetric local-lemma criterion

We use the asymmetric Lovász local lemma in the following standard form [AS16, Lemma 5.1.1]. If events (A_i) have a dependency digraph, and numbers $0 \le x_i < 1$ satisfy

$$\mathbb{P}(A_i) \le x_i \prod_{j\in\Gamma(i)} (1 - x_j) \tag{2.17}$$

for every i, then $\mathbb{P}(\bigcap_i \overline{A_i}) > 0$.

The collision events later in the proof depend on at most two independently sampled request variables. Suppose that, for every one request, the sum of the probabilities of collision events involving it is at most B. A collision event E then has

$$\sum_{E'\sim E} \mathbb{P}(E') \le 2B.$$

Set $x_E = 2\mathbb{P}(E)$. If $B \le 1/8$, then

$$\sum_{E'\sim E} x_{E'} \le 4B \le \frac{1}{2}$$

and, using $\prod_i (1 - z_i) \ge 1 - \sum_i z_i$,

$$x_E \prod_{E'\sim E} (1 - x_{E'}) \ge 2\mathbb{P}(E) \left(1 - \sum_{E'\sim E} x_{E'}\right) \ge \mathbb{P}(E).$$

Thus (2.17) holds. The factor 4 between the per-request probability bound and the dependency-neighborhood x-mass is included explicitly here.

2.6 Nagura's prime interval

Nagura proved that for every real $x > 25$ there is a prime p with

$$x < p < \frac{6}{5}x \tag{2.18}$$

[Nag52]. Consequently, if $p_- < p$ are consecutive primes and $p > 401$, then $p < (6/5)p_-$, and

$$(p-1)\sum_{r=p_-}^{p-1} \frac{1}{(r+1)(2r+1)} < \frac{p(p-p_-)}{2p_-^2} < \frac{3}{25}. \tag{2.19}$$

The finite allocation certificate is checked through 401; hence (2.19) is invoked only where the hypothesis of (2.18) is automatic.

3 The thirteen-layer lower bound

We begin with the lower bound. This is Mausberg's thirteen-layer valuation argument [Mau26], with a direct preliminary lemma that rules out all endpoints at or below $2n$. Put

$$N := \frac{n}{\log n}, \qquad \mathcal{P}_{\mathrm{sm}} := \{2, 3, 5, 7, 11, 13, 17, 19, 23\},$$

and define

$$\begin{aligned} A_{13} &:= \sum_{r=1}^{13} \frac{1}{(r+1)(2r+1)} = \frac{2014819799}{5736673800}, \\ S_{23} &:= \sum_{\ell\in\mathcal{P}_{\mathrm{sm}}} \frac{1}{\ell-1} = \frac{17927}{7920}, \\ C_0 &:= \frac{A_{13}}{S_{23}} = \frac{4029639598}{25970038185}. \end{aligned} \tag{3.1}$$

For an integer $M > n$, set

$$Q(n, M) := \frac{M!}{(n!)^2}.$$

We call M *admissible* if $Q(n, M)$ is the product of a subset of the distinct integers in $(n, M]$. This terminology is equivalent to $f(n) \le M$. Indeed,

$$\prod_{n<a\le M} a = \frac{M!}{n!}.$$

Thus, if one subset has product $n!$, its complement has product $Q(n, M)$; conversely, the complement of a subset with product $Q(n, M)$ has product $n!$.

Lemma 3.1. *For all sufficiently large n, no integer $M \le 2n$ is admissible.*

Proof. Suppose that $M \le 2n$ is admissible, and let $\mathcal{B} \subset (n, M] \cap \mathbb{Z}$ have product $Q(n, M)$. Adjoin every integer in $(M, 2n]$ and call the resulting set $\mathcal{B}'$. The factors remain distinct and

$$\prod_{b\in\mathcal{B}'} b = \frac{M!}{(n!)^2}\frac{(2n)!}{M!} = \binom{2n}{n}.$$

For $1 \le r \le 13$, let

$$\mathcal{L}_r^{(0)} := \left\{P \text{ prime} : \frac{n}{r+1} < P \le \frac{2n}{2r+1}\right\}. \tag{3.2}$$

If $P \in \mathcal{L}_r^{(0)}$, then

$$\left\lfloor \frac{n}{P} \right\rfloor = r, \qquad \left\lfloor \frac{2n}{P} \right\rfloor = 2r + 1.$$

Moreover $P > n/14$, so $P^2 > 2n$ when n is sufficiently large. Legendre's formula therefore gives

$$v_P \binom{2n}{n} = 1. \tag{3.3}$$

The layers are pairwise disjoint, since membership determines $r = \lfloor n/P \rfloor$.

Equation (3.3) forces exactly one factor of $\mathcal{B}'$ to be divisible by P. Write this factor as $Pq(P)$. Because it belongs to $(n, 2n]$, the endpoint inequalities in (3.2) imply

$$r + 1 \leq q(P) \leq 2r + 1.$$

No factor at most $2n$ contains two primes from the union of the thirteen layers: each such prime is larger than $n/14$, and the product of two of them is larger than $2n$ for large n. Consequently the forced factors $Pq(P)$ are distinct.

Every integer between 2 and 27 has a prime divisor in $\mathcal{P}_{\text{sm}}$. For each P, choose one such divisor $\ell(P) \mid q(P)$. Distinctness of the forced factors gives the incidence inequality

$$\begin{aligned} \sum_{r=1}^{13} |\mathcal{L}_r^{(0)}| &\leq \sum_{\ell \in \mathcal{P}_{\text{sm}}} \sum_{b \in \mathcal{B}'} v_\ell(b) \\ &= \sum_{\ell \in \mathcal{P}_{\text{sm}}} v_\ell \binom{2n}{n}. \end{aligned} \tag{3.4}$$

For each fixed r, the prime number theorem yields

$$|\mathcal{L}_r^{(0)}| = \left(\frac{2}{2r+1} - \frac{1}{r+1} + o(1) \right) N = \left(\frac{1}{(r+1)(2r+1)} + o(1) \right) N. \tag{3.5}$$

On the other hand, for each fixed prime ℓ,

$$v_\ell \binom{2n}{n} = \sum_{j \geq 1} \left(\left\lfloor \frac{2n}{\ell^j} \right\rfloor - 2 \left\lfloor \frac{n}{\ell^j} \right\rfloor \right) = O_\ell(\log n). \tag{3.6}$$

Indeed, every summand is either 0 or 1, and only $O_\ell(\log n)$ summands are nonzero. The left side of (3.4) is $(A_{13} + o(1))N$ by (3.5), whereas its right side is $O(\log n) = o(N)$. This contradiction proves the lemma. ∎

Lemma 3.2. *As $n \to \infty$,*

$$f(n) \geq 2n + (C_0 - o(1)) \frac{n}{\log n}.$$

Equivalently,

$$\liminf_{n \to \infty} \frac{(f(n) - 2n) \log n}{n} \geq C_0.$$

Proof. Let M be admissible. Lemma 3.1 allows us, for large n, to write

$$M = 2n + h, \qquad h > 0.$$

If $h \geq N$, the desired inequality follows from $C_0 < 1$. We may therefore assume

$$0 < h < N. \tag{3.7}$$

For $1 \le r \le 13$, define the moving layer

$$\mathcal{L}_r := \left\{ P \text{ prime} : \frac{M}{2r+2} < P \le \frac{M}{2r+1}, \quad \frac{n}{r+1} < P \le \frac{n}{r} \right\}.$$

The second pair of inequalities is redundant for all sufficiently large n, uniformly for h in (3.7). The lower inequality follows from $M > 2n$, and the upper inequality follows from

$$\frac{M}{2r+1} < \frac{n}{r} \quad \Longleftrightarrow \quad rh < n,$$

which holds for every fixed $r \le 13$ because $h < N$. The prime number theorem at the two endpoints now gives, uniformly for $0 < h < N$,

$$\begin{aligned} |\mathcal{L}_r| &= \left(\frac{2}{2r+1} - \frac{1}{r+1} + o(1) \right) N \\ &= \left(\frac{1}{(r+1)(2r+1)} + o(1) \right) N. \end{aligned} \tag{3.8}$$

For clarity, this uniformity is an immediate consequence of

$$\pi(an + O(N)) = \frac{an}{\log n} + o(N)$$

for fixed $a > 0$, uniformly when the $O(N)$-shift is in a fixed bounded range. There are only thirteen endpoint pairs.

If $P \in \mathcal{L}_r$, then

$$\left\lfloor \frac{M}{P} \right\rfloor = 2r+1, \qquad \left\lfloor \frac{n}{P} \right\rfloor = r.$$

Also $P > M/28$, and hence $P^2 > M$ for large n. Thus Legendre's formula has no higher-power terms and gives the exact identity

$$v_P(Q(n,M)) = \left\lfloor \frac{M}{P} \right\rfloor - 2 \left\lfloor \frac{n}{P} \right\rfloor = 1.$$

Let $\mathcal{B} \subset (n, M]$ be a distinct-factor representation of $Q(n,M)$. For every P in the thirteen layers, exactly one member of $\mathcal{B}$ is divisible by P; write it as $Pq(P)$. The inequalities $P \le n/r$, $Pq(P) > n$, $P > M/(2r+2)$, and $Pq(P) \le M$ imply

$$r + 1 \le q(P) \le 2r+1. \tag{3.9}$$

The layers are disjoint because $r = \lfloor n/P \rfloor$. Moreover, no factor at most M can contain two layer primes: both would exceed $M/28$, whose square exceeds M for large M. Thus the forced factors are distinct.

Choose a prime $\ell(P) \in \mathcal{P}_{\mathrm{sm}}$ dividing $q(P)$, which is possible by (3.9). The same incidence count as before gives

$$\sum_{r=1}^{13} |\mathcal{L}_r| \le \sum_{\ell \in \mathcal{P}_{\mathrm{sm}}} v_\ell(Q(n,M)). \tag{3.10}$$

Factor the quotient as

$$Q(n,M) = \binom{2n}{n} \prod_{2n < a \le 2n+h} a.$$

For a fixed prime ℓ, Legendre's formula and (3.6) yield

$$\begin{aligned} v_\ell(Q(n,M)) &= v_\ell \binom{2n}{n} + \sum_{j \ge 1} \left(\left\lfloor \frac{2n+h}{\ell^j} \right\rfloor - \left\lfloor \frac{2n}{\ell^j} \right\rfloor \right) \\ &\le \frac{h}{\ell - 1} + O_\ell(\log n). \end{aligned} \tag{3.11}$$

Here each floor difference is at most $h/\ell^j + 1$, and there are only $O_\ell(\log n)$ relevant powers.

Summing (3.8) and using (3.10)–(3.11), we obtain

$$(A_{13} + o(1))N \le hS_{23} + O(\log n).$$

Since $\log n = o(N)$, this is

$$h \ge (C_0 - o(1))N$$

by (3.1). Together with the already treated case $h \ge N$, this proves the assertion for every admissible endpoint and hence for $M = f(n)$. □

4 The certified allocation and the central anchors

For the upper bound, fix a constant

$$c > C_0 = \frac{4029639598}{25970038185},$$

and write

$$\varepsilon := c - C_0, \qquad L := \log n, \qquad N := \frac{n}{L}, \qquad h := \lceil cN \rceil, \qquad M := 2n + h. \tag{4.1}$$

We split the complement quotient as

$$C_n := \binom{2n}{n}, \qquad T_n := \prod_{1 \le j \le h} (2n + j), \qquad Q(n, M) = C_n T_n. \tag{4.2}$$

This section constructs distinct factors in $(n, 2n]$ whose product is C_n times a controlled divisor of T_n. In particular, it removes the central binomial coefficient exactly, not merely up to an asymptotic valuation error.

4.1 An exact infinite cofactor allocation

Set

$$\alpha_r := \frac{1}{(r+1)(2r+1)} \qquad (r \ge 1).$$

The following finite statement is verified entirely in rational arithmetic.

Lemma 4.1 (Finite allocation certificate)**.** *There are* 211 *positive rational numbers*

$$x^{\mathrm{fin}}_{r,q}, \qquad 1 \le r \le 200, \quad r + 1 \le q \le 2r + 1,$$

all omitted coordinates being zero, with the following properties.

1. *For every* $1 \le r \le 200$,
$$\sum_{q=r+1}^{2r+1} x^{\mathrm{fin}}_{r,q} = \alpha_r.$$

2. *If*
$$\lambda^{\mathrm{fin}}_\ell := \sum_{r=1}^{200} \sum_q v_\ell(q) x^{\mathrm{fin}}_{r,q},$$
then
$$\lambda^{\mathrm{fin}}_\ell \le \frac{C_0}{\ell - 1}$$
for every prime $\ell \le 401$.

3. If $201 < p \le 401$ *is prime and* p_- *denotes the prime immediately preceding* p*, then*

$$\lambda_p^{\mathrm{fin}} + \sum_{r=\max(201,p_-)}^{p-1} \alpha_r \le \frac{C_0}{p-1}.$$

The complete 211-entry rational array and its exact checker are contained in the single companion file `numerical_verifier.py`, distributed with this paper. The checker verifies positivity and the permitted cofactor range, adds every row in exact rational arithmetic, factors every cofactor, and checks the two families of capacity inequalities in Lemma 4.1. Thus the file is both an explicit certificate for the existential statement above and a reproducible verification of every finite calculation used here. No floating-point calculation enters the verification; see Appendix A for its precise scope.

Lemma 4.2 (Infinite cofactor allocation)**.** *There are nonnegative rational numbers* $x_{r,q}$*, indexed by*

$$r \ge 1, \qquad r+1 \le q \le 2r+1,$$

such that

$$\sum_{q=r+1}^{2r+1} x_{r,q} = \alpha_r \qquad\qquad (r \ge 1), \tag{4.3}$$

$$\sum_{r\ge 1}\sum_{q=r+1}^{2r+1} v_\ell(q) x_{r,q} \le \frac{C_0}{\ell-1} \qquad\qquad (\ell \textit{ prime}). \tag{4.4}$$

Proof. For $r \le 200$, use the array in Lemma 4.1. For $r \ge 201$, let $p(r)$ be the least prime strictly larger than r, and put

$$x_{r,p(r)} := \alpha_r, \qquad x_{r,q} := 0 \quad (q \ne p(r)).$$

Bertrand's postulate gives $r < p(r) < 2r$, so $p(r)$ is an allowed cofactor. The row identities (4.3) follow.

It remains to check the prime capacities. A tail row contributes only to the prime $p(r)$. If p and p_- are consecutive primes, the indices for which the least-prime rule routes to p are precisely

$$p_- \le r \le p-1.$$

The actual tail contribution is the intersection of this block with $r \ge 201$. For $p \le 200$, there is no tail contribution, while for $200 < p \le 401$, the finite-plus-tail inequality is exactly the third check in Lemma 4.1.

Now let $p > 401$. Every finite cofactor is at most 401, so its p-load is zero. Nagura's theorem [Nag52], applied to $p_- > 25$, gives $p < 6p_-/5$. Hence $p - p_- < p_-/5$, and

$$\begin{aligned}(p-1)\sum_{r=p_-}^{p-1} \alpha_r &< p\,\frac{p-p_-}{2p_-^2} \\ &< \frac{3}{25} < C_0.\end{aligned} \tag{4.5}$$

The first strict inequality uses $(r+1)(2r+1) > 2p_-^2$ throughout the block, and the last one follows directly from the rational value of C_0. Dividing (4.5) by $p-1$ proves (4.4) for every remaining prime. ∎

4.2 Exact realization of a fixed prefix

For $r \geq 1$, define the central carry row

$$\mathcal{C}_r(n) := \left\{ P \text{ prime} : \frac{n}{r+1} < P \leq \frac{2n}{2r+1} \right\}. \tag{4.6}$$

If $P \in \mathcal{C}_r(n)$, then

$$\left\lfloor \frac{n}{P} \right\rfloor = r, \qquad \left\lfloor \frac{2n}{P} \right\rfloor = 2r+1.$$

For fixed r and sufficiently large n, one also has $P^2 > 2n$, and therefore

$$v_P(C_n) = 1.$$

Every integer q in the allocation range obeys

$$r+1 \leq q \leq 2r+1 \quad \Longrightarrow \quad n < Pq \leq 2n. \tag{4.7}$$

The prime number theorem gives, for each fixed r,

$$|\mathcal{C}_r(n)| = (\alpha_r + o(1))N. \tag{4.8}$$

Fix an integer $R \geq 200$, independently of n. For each $r \leq R$, partition $\mathcal{C}_r(n)$ among the positive coordinates $x_{r,q}$. More precisely, round $x_{r,q}N$ to integers for all but one positive coordinate and assign all remaining primes in the row to the last coordinate. Because the number of coordinates is fixed and (4.8) holds, the resulting parts $\mathcal{C}_{r,q}(n)$ satisfy

$$|\mathcal{C}_{r,q}(n)| = x_{r,q}N + o(N) \tag{4.9}$$

for every positive coordinate; in particular, the last part is nonnegative for all sufficiently large n. The parts exhaust the row exactly.

Use the factors

$$\mathcal{H}_R^{\text{pre}} := \{Pq : P \in \mathcal{C}_{r,q}(n),\ 1 \leq r \leq R\}. \tag{4.10}$$

They lie in $(n, 2n]$ by (4.7). They are also distinct. Indeed, for fixed R and large n,

$$P > \frac{n}{R+1} > 2R+1 \geq q.$$

Thus P is the unique prime divisor of Pq larger than $2R+1$; equality of two prefix anchors forces equality of both P and q. This also shows that repeated use of the same cofactor in different rows cannot create a collision.

Define the cofactor product

$$D_R^{\text{pre}} := \prod_{r=1}^{R} \prod_{q} q^{|\mathcal{C}_{r,q}(n)|}.$$

For every prime ℓ, the finite realization (4.9) and the full capacity inequality (4.4) give

$$v_\ell(D_R^{\text{pre}}) \leq \left(\frac{C_0}{\ell - 1} + o_R(1) \right) N. \tag{4.11}$$

All prime divisors of D_R^{pre} are at most $2R+1$.

There is also a row $r = 0$. Every prime P with

$$n < P \leq 2n \tag{4.12}$$

occurs once in C_n, and we use P itself as an anchor. These singletons have cofactor 1, consume no tail valuation, and cannot collide with the composite factors in (4.10).

4.3 Promotion of all residual central valuations

Put

$$X_R^{\mathrm{anc}} := \frac{n}{R+1}.$$

Once R is fixed, take n sufficiently large that

$$X_R^{\mathrm{anc}} > 2R+1, \qquad (X_R^{\mathrm{anc}})^2 > 2n.$$

Every prime $P > X_R^{\mathrm{anc}}$ occurring in C_n then occurs to the first power and is accounted for by exactly one of the rows (4.6) with $1 \le r \le R$, or by the row-zero set (4.12). Indeed, $\lfloor n/P \rfloor \le R$, and the equality

$$\left\lfloor \frac{2n}{P} \right\rfloor - 2\left\lfloor \frac{n}{P} \right\rfloor = 1$$

is precisely the membership condition for the corresponding carry row.

It remains to account for primes $p \le X_R^{\mathrm{anc}}$. Put

$$e_p := v_p(C_n), \qquad B_p := p^{e_p}.$$

Every summand in Legendre's formula

$$e_p = \sum_{j \ge 1} \left(\left\lfloor \frac{2n}{p^j} \right\rfloor - 2\left\lfloor \frac{n}{p^j} \right\rfloor \right)$$

is 0 or 1. Consequently

$$e_p \le \lfloor \log_p(2n) \rfloor, \qquad B_p \le 2n. \tag{4.13}$$

For every $p \le X_R^{\mathrm{anc}}$ with $e_p > 0$, let k_p be the least nonnegative integer for which

$$n < 2^{k_p} B_p. \tag{4.14}$$

Minimality and (4.13) imply

$$n < 2^{k_p} B_p \le 2n.$$

We therefore use the single promoted factor

$$A_p := 2^{k_p} p^{e_p}. \tag{4.15}$$

The promoted factors are mutually distinct. Factors belonging to two different odd base primes are distinguished by unique factorization, and the factor with base prime 2 is a pure power of two. They also cannot collide with an earlier anchor: every prime divisor of a promoted factor is at most X_R^{anc}, whereas a prefix anchor has a unique prime divisor above X_R^{anc}, and a row-zero singleton is itself a prime above n.

The additional power of two has small total cost. Define

$$K_R(n) := \sum_{\substack{p \le X_R^{\mathrm{anc}} \\ e_p > 0}} k_p.$$

Since $B_p \ge p$, minimality in (4.14) gives

$$k_p \le 1 + \log_2 \frac{n}{p}. \tag{4.16}$$

For primes $p \le \sqrt{n}$, the total contribution is

$$O(\pi(\sqrt{n}) \log n) = O(\sqrt{n}) = o(N).$$

For $p > \sqrt{n}$, split the primes into dyadic blocks

$$\frac{n}{2^{j+1}} < p \leq \frac{n}{2^j}.$$

Only indices

$$j \geq \lfloor \log_2(R+1) \rfloor - O(1), \qquad 2^j \leq \sqrt{n},$$

occur. The standard Chebyshev upper bound for primes, uniformly in this range, gives

$$\pi(n/2^j) \ll \frac{n}{2^j \log(n/2^j)} \ll \frac{N}{2^j}.$$

Together with (4.16), this yields

$$\begin{aligned} K_R(n) &\ll N \sum_{j \geq \log_2(R+1) - O(1)} \frac{j+1}{2^j} + o(N) \\ &\ll N \frac{1 + \log R}{R} + o(N), \end{aligned}$$

with an absolute implied constant. In particular,

$$\lim_{R \to \infty} \limsup_{n \to \infty} \frac{K_R(n)}{N} = 0. \tag{4.17}$$

4.4 The exact anchor product and its tail reserve

Let $\mathcal{H}_R$ be the union of the prefix anchors (4.10), the row-zero singletons (4.12), and the promoted factors (4.15). The preceding arguments show that these are distinct integers in $(n, 2n]$. Every prime valuation of C_n appears exactly once in their product: the valuations above X_R^{anc} occur in their carry or row-zero anchors, and each valuation at or below X_R^{anc} occurs in its complete block p^{e_p}. Therefore, if

$$D_R := D_R^{\text{pre}} \, 2^{K_R(n)}, \tag{4.18}$$

then the following is an exact integer identity:

$$\prod_{a \in \mathcal{H}_R} a = C_n D_R. \tag{4.19}$$

Moreover D_R is supported on the fixed prime set

$$F_R := \{\ell \text{ prime} : \ell \leq 2R + 1\}.$$

We now choose $R = R(c)$. By (4.17), we may fix $R \geq 201$ so large that

$$K_R(n) \leq \frac{\varepsilon}{4} N$$

for every sufficiently large n. This choice is made before n tends to infinity. For each fixed prime ℓ, Legendre's formula over the tail interval gives

$$\begin{aligned} v_\ell(T_n) &= \sum_{j \geq 1} \left(\left\lfloor \frac{2n+h}{\ell^j} \right\rfloor - \left\lfloor \frac{2n}{\ell^j} \right\rfloor \right) \\ &= \frac{h}{\ell - 1} + O_\ell(\log n) = \left(\frac{c}{\ell - 1} + o(1) \right) N. \end{aligned} \tag{4.20}$$

For an odd prime $\ell \in F_R$, (4.11) and (4.18) give

$$v_\ell(D_R) \leq \left(\frac{C_0}{\ell - 1} + o(1) \right) N.$$

At $\ell = 2$, the promotion cost adds at most $\varepsilon N/4$, so

$$v_2(D_R) \le \left(C_0 + \frac{\varepsilon}{4} + o(1)\right) N.$$

Because F_R is fixed, these inequalities and (4.20) hold simultaneously at all of its primes. After increasing n, if necessary, they imply

$$v_\ell(T_n) - v_\ell(D_R) \ge \frac{\varepsilon}{3(\ell-1)} N \qquad (\ell \in F_R). \tag{4.21}$$

Outside F_R, the divisor D_R has zero valuation. Thus

$$D_R \mid T_n. \tag{4.22}$$

We record the conclusion in the form used later.

Lemma 4.3 (Full central-anchor lemma). *Fix $c > C_0$, and let h, M, C_n, T_n be as in (4.1)–(4.2). There is a finite prime set $F_{\text{anc}} = F_{\text{anc}}(c)$ such that, for all sufficiently large n, one can find an F_{anc}-supported divisor $D \mid T_n$ and a set $\mathcal{H}$ of distinct integers in $(n, 2n]$ satisfying*

$$\prod_{a \in \mathcal{H}} a = C_n D.$$

In addition,

$$v_\ell(T_n) - v_\ell(D) \ge \frac{c - C_0}{3(\ell-1)} N \qquad (\ell \in F_{\text{anc}}). \tag{4.23}$$

Consequently, removal of these anchors leaves the exact quotient

$$\frac{Q(n, M)}{\prod_{a \in \mathcal{H}} a} = \frac{T_n}{D}.$$

Proof. Take $R = R(c)$ as above, set $F_{\text{anc}} := F_R$, $D := D_R$, and $\mathcal{H} := \mathcal{H}_R$. Equations (4.19), (4.21), and (4.22) give all assertions. ☐

We finally note an endpoint detail that will be useful when the residual tail is realized. A signature prime $P > X_R^{\text{anc}}$ used in a carry or row-zero anchor can occur once in C_n and also occur in one or more tail terms. The anchor construction uses exactly its central occurrence. Since D is supported on primes at most $2R + 1 < X_R^{\text{anc}}$, all of these additional P-occurrences remain in T_n/D. This causes no numerical collision: every anchor is at most $2n$, whereas every actual tail term is larger than $2n$. In particular, if $n < P \le n + h/2$, the row-zero anchor is the singleton P, while $2P$, if it is used later, is a distinct tail factor. Primes in $(2n, M]$ occur only in the tail and are untouched by the anchor construction.

5 A precharged two-sided universal bank

We reserve now the switches that will absorb the discrepancy produced by integral rounding. This is done before exceptional tail factors are retained and before a fractional selector is solved: the chosen base state of every switch is part of the charged residual product.

Put

$$\theta := \frac{2}{9}, \qquad y := n^\theta, \qquad d_n := \lceil \log_2(3n) \rceil, \qquad \beta_p := 4 d_n \left\lceil \frac{\log(3n)}{\log p} \right\rceil \qquad (p \le y).$$

Lemma 10.1 will show that retaining every integer rough-row sum while rounding creates an integral error e, supported on $p \le y$, with $|e_p| \le \beta_p$. The elementary PNT estimates

$$\sum_{p \le y} \frac{1}{\log p} \ll \frac{y}{(\log y)^2}, \qquad \pi(y) \ll \frac{y}{\log y}$$

give

$$\sum_{p\le y}\beta_p = O(y), \tag{5.1}$$

$$\sum_{p\le y}\beta_p \log p = O(yL).$$

5.1 Complete signatures and path states

For every legal lower or upper factor define its complete rough signature

$$S(a) := (v_P(a))_{P>y} \qquad (n < a \le M).$$

Equivalently, the integer label used in Section 6 is $R_y(a) = \prod_{P>y} P^{S_P(a)}$; the vector and integer conventions partition the factors into exactly the same rows. For a finite set G of legal factors, put

$$m_S(G) := \#\{a \in G : S(a) = S\}, \qquad v(G) := v\left(\prod_{a\in G} a\right).$$

A switch g has two states G_g^0, G_g^1 and change $\Delta_g = v(G_g^1) - v(G_g^0)$. We require

$$m_S(G_g^0) = m_S(G_g^1) \qquad \text{for every complete signature } S. \tag{5.2}$$

Our switches are paths of component edges. Each component edge has one factor in each state, and those factors have the same complete signature; different components use different marker primes. Hence a complete path state has one token in each of its component marker rows, and its two full states have identical signature counts row by row. In particular, a full path has one token per component row, not one token in total.

Denote the eventual collection of path switches by $\mathcal{B}$. It is *two-sided β-universal* if every integer vector

$$z = (z_p)_{p\le y}, \qquad |z_p| \le \beta_p, \tag{5.3}$$

is the exact valuation change of a collision-free collection of toggles. We now construct such a bank while backing every one of its rough tokens by a distinct actual tail occurrence.

5.2 The literal geometric descent

Use the grid

$$Q_j := 4(4/3)^j \qquad (j \ge 0).$$

Every integer core $q > 4$ lies in a unique cell $(Q, 4Q/3]$. At a nonbottom cell choose

$$q \in (Q, 4Q/3], \qquad 3Q/4 < b < q, \qquad P \in I_Q := \left(\frac{4n}{3Q}, \frac{3n}{2Q}\right].$$

Then

$$n < Pb < Pq \le 2n,$$

so $Pq \leftrightarrow Pb$ is a legal lower switch of ratio b/q.

Here is a literal descent which avoids the power-of-two cores of the promoted anchors. If $q \ge 6$ is not a power of two and $Q > 20$, set $b_0 = \lceil 4q/5 \rceil$ and

$$b := \begin{cases} b_0, & b_0 \text{ is not a power of two,} \\ b_0 - 1, & b_0 \text{ is a power of two.} \end{cases}$$

Then b is a non-power and

$$b > \frac{4}{5}Q - 1 > \frac{3}{4}Q, \qquad 5 \le b < q.$$

For the non-power cores in cells with $Q \le 20$, use

q	6	7	9	10	11	12	13	14	15	17	18	19	20	21	22
b	5	6	7	9	9	10	11	12	12	14	15	15	15	17	18.

Every displayed pair satisfies the strict cell inequalities. A source prime $p \ge 5$ is a non-power, so induction keeps every intermediate core a non-power. The core decreases, and above the table it is at most $4q/5 + 1$; the path therefore reaches 5 in $O(\log p)$ steps. We stop there. Moreover, when $Q > 20$ one has $b < 17q/20$; hence two consecutive moves whose source cores remained in the same cell would reduce a core at most $4Q/3$ below $(17/20)^2(4Q/3) < Q$. Thus a path uses at most two edges at every large scale, while the finite table gives an absolute bound at the five small scales. In particular it uses $O(1)$ edges at each geometric scale. The strict ordinary descent is not asserted to reach 2.

For a nonbottom scale $Q \le y$, define the donor multiplicity

$$s_P := \#\left\{u : \frac{7Q}{5} \le u \le \frac{29Q}{20},\ P^+(u) \le y,\ 2n < Pu \le M\right\}. \tag{5.4}$$

When $Q > 20$, there are $\asymp Q$ eligible y-smooth u's. Indeed, $u < 2y$, so an integer here having a prime factor above y must itself be such a prime; this discards only $O(Q/\log Q)$ values. For fixed eligible u, the interval

$$2n/u < P \le M/u$$

lies in a fixed shrunken part of I_Q. The uniform PNT (2.8) gives

$$\#\{P : 2n/u < P \le M/u\} \asymp_c \frac{N}{QL}. \tag{5.5}$$

Indeed its endpoints have size n/Q, with $\log(n/Q) \ge (1-\theta)L$, and its length is $h/u \asymp_c N/Q$. The cumulative-PNT error at the two moving endpoints is $O((n/Q)e^{-c_0\sqrt{L}}) = o(N/(QL))$, uniformly for $Q \le y$. Thus (5.5) does not invoke a prime theorem for arbitrary short intervals. At the five nonbottom scales $Q_j \le 20$, use respectively the fixed donor cofactors

j	1	2	3	4	5
u	8	10	13	17	23.

Each lies between $4Q_j/3$ and $3Q_j/2$, and the same PNT estimate applies.

Double-counting (P, u) gives

$$C_Q := \sum_{P \in I_Q} s_P \asymp_c \frac{N}{L}. \tag{5.6}$$

A fixed $P \in I_Q$ admits only $O(h/P + 1) = O_c(Q/L + 1)$ donor cofactors. Hence

$$K_Q := \#\{P \in I_Q : s_P > 0\} \asymp_c \frac{N}{\max(Q, L)}. \tag{5.7}$$

The upper bound follows from either the length of I_Q or (5.6); division by the maximum multiplicity gives the lower bound.

Choose one donor for each of the K_Q eligible marker primes, and split this set of marker–donor pairs into two parts whose sizes differ by at most one. These are the two ordinary orientation pools. Thus each orientation separately retains $\asymp K_Q$ distinct signatures; the occurrence bound (5.6) remains available as a separate verification of donor load.

Successive marker intervals are disjoint: if $Q' = 4Q/3$, then

$$I_{Q'} = \left(\frac{n}{Q}, \frac{9n}{8Q}\right], \qquad I_Q = \left(\frac{4n}{3Q}, \frac{3n}{2Q}\right],$$

with a fixed gap. Moreover

$$P \gg n/y > y, \qquad P^2 > M. \tag{5.8}$$

Thus the donor and both endpoints have the same complete signature, namely the single marker P with exponent one.

5.3 Four bottom pools and the signed unit lattice

Continue the descent from 5 through the four pairwise disjoint pools

move	marker interval	states	tail donor
$5 \to 4$	$(n/3, M/6]$	$5P \leftrightarrow 4P$	$6P$
$4 \to 3$	$(2n/5, M/5]$	$4P \leftrightarrow 3P$	$5P$
$3 \to 2$	$(2n/3, M/3]$	$3P \leftrightarrow 2P$	$3P$
$2 \to 1$	$(n/2, M/4]$	$4P \leftrightarrow 2P$	$4P.$

(5.9)

All state factors are in $(n, M]$. In the last two lines the upper state is itself the tail donor; it is one occurrence, not a retained donor plus a second bank factor. The unsplit prime capacities are

$$\left(\frac{c}{6} + o(1)\right)\frac{N}{L}, \quad \left(\frac{c}{5} + o(1)\right)\frac{N}{L}, \quad \left(\frac{c}{3} + o(1)\right)\frac{N}{L}, \quad \left(\frac{c}{4} + o(1)\right)\frac{N}{L}.$$

These are differences of the cumulative PNT at endpoints of size $\asymp n$: the four interval lengths are $h/6, h/5, h/3, h/4$, and the absolute PNT remainder $O(ne^{-c_0\sqrt{L}})$ is $o(N/L)$. Split each interval into two fixed positive subintervals for the two orientations.

The word "split" means a split into two subintervals whose lengths are fixed positive proportions of the original interval. Consequently, for constants $c_i = c_i(c) > 0$, each of the eight oriented bottom pools contains

$$(c_i + o(1))\frac{N}{L} \quad \text{marker primes.} \tag{5.10}$$

This is much larger than the $O(y)$ requests made to any one pool.

For $p \geq 5$, concatenate the ordinary $p \to 5$ path and the four bottom moves. In the downward orientation its ratio and valuation change are

$$\frac{4}{p}\frac{3}{4}\frac{2}{3}\frac{1}{2} = \frac{1}{p},$$
$$(2\mathbf{e}_2 - \mathbf{e}_p) + (\mathbf{e}_3 - 2\mathbf{e}_2) + (\mathbf{e}_2 - \mathbf{e}_3) - \mathbf{e}_2 = -\mathbf{e}_p. \tag{5.11}$$

For $p = 3$, use only $3 \to 2 \to 1$, and for $p = 2$, only the terminal move. These yield $-\mathbf{e}_3$ and $-\mathbf{e}_2$. Reverse every base orientation in a disjoint copy to obtain the positive unit vector. Reserve β_p copies of each orientation for every $p \leq y$.

There are $O(y)$ paths by (5.1). A path uses only $O(1)$ edges at any geometric scale. At the worst scale $Q = y$,

$$\frac{O(y)}{K_Q} \ll \frac{y^2}{N} = o(1), \qquad \frac{O(y)}{C_Q} \ll \frac{yL}{N} = o(1). \tag{5.12}$$

Thus every ordinary component can receive its own marker prime and donor. Each bottom pool receives $O(y)$ requests, and its separate utilization is

$$O(y)/(N/L) = O(yL/N) = o(1). \tag{5.13}$$

The total number of component factors is

$$O\left(\sum_{p\le y} \beta_p \log p\right) = O(yL) = o(N). \tag{5.14}$$

Lemma 5.1 (Disjoint donor assignment)**.** *For all sufficiently large n, every component in the fully reserved bank can be assigned a pair (P, Pu), where P is its marker and $Pu \in E$ is its backing donor, so that the following hold simultaneously.*

1. *Distinct components have distinct markers and distinct donor occurrences.*
2. *The two orientations at a fixed scale use disjoint pools; pools at different ordinary scales, and the four bottom pools, are mutually disjoint.*
3. *The donor and both states of the component have the same complete rough signature.*
4. *If a bottom donor is itself one of the two states, it denotes that one occurrence, not an additional copy of the same integer.*

Thus assignment of the donors is an injection, rather than only a comparison of their total number with the number of requests.

Proof. At an ordinary scale Q, first choose one donor for every eligible marker, as above, and split the resulting marker–donor pairs into the two orientation pools. Let $m_Q^{\pm}$ be the numbers of component requests of the two orientations. A path has $O(1)$ components at a given scale, by the literal descent estimate preceding (5.12), and (5.1) gives $m_Q^{\pm} = O(y)$. On the other hand, each orientation pool has $\asymp K_Q$ pairs. Uniformly for $Q \le y$,

$$\frac{m_Q^{\pm}}{K_Q} \ll \frac{y \max(Q, L)}{N} \le \frac{y^2 + yL}{N} = o(1)$$

by (5.7). A greedy assignment therefore uses a fresh pair for every request. Successive ordinary marker intervals are disjoint by the calculation following (5.7), so assignments made at different scales cannot meet.

For a bottom pool, (5.10) supplies $(c_i + o(1))N/L$ marker–donor occurrences, whereas the request count is $O(y) = o(N/L)$. The same greedy construction works separately in each oriented subpool. The four displayed marker intervals are disjoint for large n, and the last ordinary marker interval lies below the first bottom interval, so these choices do not meet any ordinary choice.

Finally, at an ordinary scale the three cofactors q, b, u are y-smooth and $P > y$, $P^2 > M$, by (5.8); the same statement is immediate from the four fixed bottom rows. Hence every factor in one component has complete signature P with exponent one. In the last two bottom rows the donor and one state coincide by definition, and the construction counts that single occurrence only once. This proves all assertions. □

Given z in (5.3), toggle $|z_p|$ copies having its sign. Equation (5.11), with its $p = 2, 3$ truncations, gives total change exactly z. Every such plan is a subplan of the fully reserved plan, so (5.12)–(5.13) give uniform $o(1)$ congestion throughout the box.

5.4 Anchor modifications and guards

Choose all bank markers before freezing the integer anchor partitions of Section 4. The following lemma records every point at which the resulting modification is used later.

Lemma 5.2 (Collision-free modification of the anchors)**.** *The integer partitions in the fixed prefix can be chosen so that no bank state is an anchor. More precisely, there are modified anchors $\mathcal{H}'$ and an F_{anc}-supported integer D' such that*

1. *every forced bottom marker belongs to the asserted central carry row, and every replacement cofactor is legal in that row;*
2. *the forced row-1 and row-2 markers use only $o(N)$ of positive certificate cells of size $\asymp N$;*
3. *the anchors in $\mathcal{H}'$ are mutually distinct and avoid both incident states of every bank component;*
4. $$\prod_{a\in\mathcal{H}'} a = C_n D', \qquad D' \mid T_n;$$
5. *for every $f \in F_{\text{anc}}$, the change from the original cofactor product is $o(N)$, so a fixed positive linear part of (4.23) remains.*

These statements remain valid when a legal replacement cofactor was a zero coordinate of the rational allocation: only $O(y) = o(N)$ such occurrences are inserted.

Proof. For a $5 \to 4$ marker $P \in (n/3, M/6]$, one has, for all sufficiently large n,

$$\lfloor n/P \rfloor = 2, \qquad \lfloor 2n/P \rfloor = 5, \qquad v_P(C_n) = 1.$$

Thus it lies in the row-2 carry interval, and we force its anchor to be $3P$. The relevant certificate coordinate is strictly positive:

$$x_{2,3} = \frac{3432239399}{51940076370} > 0.$$

Likewise a terminal marker $P \in (n/2, M/4]$ belongs to the row-1 carry interval and is assigned the anchor $3P$, where

$$x_{1,3} = \frac{597400199}{51940076370} > 0.$$

Each of these two certificate cells contains $x_{r,3}N + o(N) \asymp N$ marker primes by (4.9). The number forced into either cell is $O(y) = o(N)$, by (5.1); hence the integer partition can accommodate every forced marker. The terminal bank switch is still only $4P \leftrightarrow 2P$: the anchor $3P$ uses the central P-occurrence, while either bank state uses its separately backed tail occurrence.

We record explicitly why the fixed prefix meets only $O(y)$ bank components, although the entire bank has $O(yL)$ components. If an ordinary component at scale Q meets a fixed-prefix anchor, equality of the two integers forces equality of their complete rough signatures and hence forces the component marker P to be the anchor prime: for fixed R and large n, its prime is $> n/(R+1) > y$, whereas its cofactor is at most $2R + 1 < y$. A fixed-prefix anchor has $P > n/(R+1)$, while $P \in I_Q \subset (4n/(3Q), 3n/(2Q)]$. Consequently

$$Q < \frac{3(R+1)}{2}.$$

There are only $O_R(1)$ geometric scales satisfying this inequality. By the literal descent estimate preceding (5.12), each of the $O(y)$ paths has $O(1)$ components at any one scale. (The stronger assertion "at most one" is neither needed nor generally true for the finite descent table.) Hence ordinary components meeting the prefix number $O_R(y)$. The four bottom moves contribute at most another $O(y)$. Thus

$$\#\{\text{bank components meeting a fixed-prefix anchor}\} = O_R(y) = O(y). \tag{5.15}$$

At any other fixed prefix row $r \geq 2$, the legal cofactors $r+1, \ldots, 2r+1$ contain at least three integers. At most two are the incident component cores, so choose a third legal cofactor. It need not have had positive limiting allocation weight: by (5.15), adding only $O(y) = o(N)$ occurrences changes every fixed capacity by $o(N)$. Beyond the fixed prefix, a promoted nonsmooth anchor has power-of-two smooth core, whereas the ordinary descent uses non-power cores. The only unavoidable core 4 is covered by the forced row-2 assignment; the $4 \to 3$ and $3 \to 2$ bottom pools have no competing central occurrence.

Distinctness is now literal. A prefix anchor has its unique marker prime $P > n/(R+1) > 2R+1$; different components use different markers, and for a fixed marker we selected a cofactor different from both state cores. Promoted anchors were already mutually distinct, and the power-of-two-core observation separates them from ordinary states.

For each changed occurrence let a_i and b_i be its old and new cofactors. If D is written as the product of its individual cofactor occurrences, deletion of the selected occurrences and insertion of their replacements gives the integer

$$D' := D \prod_i \frac{b_i}{a_i}.$$

This notation does not assert $a_i \mid b_i$: integrality follows because the displayed denominators cancel the corresponding factors in that occurrence-level product. Every b_i lies in a fixed-prefix legal range, so all its prime divisors belong to F_{anc}. Moreover, by (5.15), for each $f \in F_{\mathrm{anc}}$,

$$|v_f(D') - v_f(D)| = O_{F_{\mathrm{anc}}}(y) = o(N). \tag{5.16}$$

Replacing the corresponding anchors changes their product by exactly the same ratio, and hence

$$D' \mid T_n, \qquad \prod_{a \in \mathcal{H}'} a = C_n D'. \tag{5.17}$$

Here divisibility and the surviving positive linear reserve follow from (4.23) and the $o(N)$ loss in (5.16). This proves the lemma. $\square$

Use a different marker prime at every component, across paths, scales, and orientations. Because every cofactor is $O(y) < P$, equality of two bank factors forces equality of their marker and then of their core. Donors have y-smooth cofactor and the same unique marker, and each actual donor is used once. Donors exceed $2n$, so they cannot equal a lower endpoint or anchor. Remove both possible states from later flexible lists, and guard them from every fixed residual factor. These choices make anchors, bank states, donors, and all later fixed or flexible factors jointly collision-free.

5.5 Donor-backed row tokens and the base charge

For each ordinary component choose one donor counted by (5.4); for each bottom component use the donor in (5.9). Remove it before a retained list or row quota is frozen and replace its signature token by the chosen base-state factor. If the donor is an upper state, the upper orientation keeps that occurrence and the lower orientation replaces it; it is never used twice.

For each component the donor and either endpoint have the same complete signature and contribute one factor to that row. The adjusted quota is therefore nonnegative and invariant under a toggle. For a full path this replacement is made separately in every component marker row, which proves (5.2) for the full states.

Fix every zero orientation and put

$$G^0_{\mathrm{bank}} := \bigcup_{g\in\mathcal{B}} G^0_g, \qquad B^0_{\mathrm{bank}} := \prod_{a\in G^0_{\mathrm{bank}}} a.$$

Unused copies remain in this charged state. Uniformly for $p \le y$, (5.14) gives

$$v_p(B^0_{\mathrm{bank}}) \ll yL\frac{L}{\log p} = o\left(\frac{N}{pL}\right). \tag{5.18}$$

Also, uniformly in this range,

$$v_p(T_n) = \frac{h}{p-1} + O\left(\frac{L}{\log p}\right) = \left(\frac{c}{p-1} + o(1)\right)N.$$

For $p \in F_{\mathrm{anc}}$, use the reserve left by (5.16); outside that set, $v_p(D') = 0$. For $P > y$, the one-to-one donor assignment injects the complete rough valuation of every base token into distinct factors of T_n. These coordinatewise statements prove the simultaneous preliminary charge

$$D'B^0_{\mathrm{bank}} \mid T_n.$$

This concerns only the bank. Section 6 checks the combined product with the retained exceptional factors before solving the fractional equations; those and the bank states exhaust the fixed residual token classes.

Lemma 5.3 (Precharged universal bank)**.** *There are modified anchors $\mathcal{H}'$, a divisor D', and a jointly guarded bank such that:*

1. *$\prod_{a\in\mathcal{H}'} a = C_nD'$ and $D'B^0_{\mathrm{bank}} \mid T_n$;*
2. *both states of every path have one token in each component marker row and satisfy (5.2);*
3. *every integer z supported on $p \le y$, with $|z_p| \le \beta_p$, is an exact collision-free bank change; and*
4. *the fully reserved bank has $o(1)$ utilization at each ordinary scale and in each of the four oriented bottom pools, and has $O(yL) = o(N)$ component factors.*

Every donor and both states have been removed or guarded before later retained lists and row quotas are defined.

6 The guarded rough-signature selector

We now construct the fractional point which fixes every nontrivial large-prime signature. Throughout this section

$$L = \log n, \qquad N = \frac{n}{L}, \qquad \theta = \frac{2}{9}, \qquad y = n^\theta,$$

and $h, M, T = T_n$ are as in (4.1) and (4.2). The divisor supplied by the anchor and bank construction will be denoted by D'. It is supported on the fixed set F_{anc}, and the replacement of D by D' preserves a fixed positive fraction of the reserve in (4.23).

We fix the order of all choices used in this and the later correction stages. After c, the prefix length $R = R(c)$, the set F_{anc}, and the modified anchor construction are fixed. Next choose constants

$$1 < r_0 < \frac{3}{2}, \qquad 1 < \rho, \qquad \rho^3 < r_0,$$

for the common-list construction of Section 9. Then choose a sufficiently large nonprime real number $W > \max F_{\mathrm{anc}}$, before any tilt box or fine mesh is chosen, and put

$$F_{\mathrm{hd}} := \{p : p \le W\}, \qquad P_{\mathrm{hd}} := \prod_{p \in F_{\mathrm{hd}}} p, \qquad \delta_{\mathrm{hd}} := \frac{\varphi(P_{\mathrm{hd}})}{P_{\mathrm{hd}}}.$$

Thus $F_{\mathrm{anc}} \subset F_{\mathrm{hd}}$, but these two sets have different roles: only the former supports D', while the primes in $F_{\mathrm{hd}} \setminus F_{\mathrm{anc}}$ retain their full tail supply. All constants depending on W are fixed before n tends to infinity.

For precision, choose constants $\gamma > 0$, $A_{\mathrm{exc}} < \infty$, and $C_{\mathrm{sv}} < \infty$, depending only on parameters already fixed, with the following meanings. For all sufficiently large n, the valuation available after the anchor charge is at least $\gamma N/p$, uniformly for every $p \le y$: at $p \in F_{\mathrm{anc}}$ this is the surviving reserve from Lemma 5.2, and otherwise it follows from $v_p(T_n) = (c + o(1))N/(p-1)$. The constant A_{exc} is an admissible constant in the nonnegative estimate (6.47) below, and C_{sv} is an admissible constant in the exceptional-list estimate (9.6). Their existence is proved without using δ_* except through the displayed linear factor.

We now choose δ_* once, before defining any cutoff-dependent set, so that

$$0 < \delta_* < \min\left\{\frac{1}{18}, \frac{\theta\gamma}{4A_{\mathrm{exc}}}, \frac{\theta\delta_{\mathrm{hd}}(2-r_0)}{4C_{\mathrm{sv}}}\right\}, \qquad X_0 = n^{\delta_*}. \tag{6.1}$$

It is never decreased later. After the remaining finite-dimensional parameters (including a permitted fixed mesh and its tilt box) have been chosen in the order specified in Section 8, there is a threshold n_0 depending on all those fixed choices. No uniformity of n_0 over a varying family of meshes is asserted. Throughout this section, $o_n(1)$ means that all parameters just listed are fixed first; an $O_W(1)$ constant may depend on the already fixed head cutoff but not on n, the later tilt-box radius, or a later mesh. Equivalently, if $\mathfrak{P}(\eta_0)$ denotes the permitted fixed meshes of Section 8, including their subsequently fixed finite-dimensional data, the quantifiers used for the upper construction have the form

$$\begin{gathered}\forall c > C_0 \ \exists (R, F_{\mathrm{anc}}, r_0, \rho, W, \delta_*) \ \exists (\eta_0, C_{\mathrm{tan}}), \\ \forall \mathcal{P} \in \mathfrak{P}(\eta_0) \ \exists n_0(\mathcal{P}) \ \forall n \ge n_0(\mathcal{P}) : \\ \text{all stated construction conclusions hold.}\end{gathered} \tag{6.2}$$

The bridge proves that η_0 and C_{tan} are independent of the particular permitted fine mesh; only n_0 is allowed to depend on that mesh. Display (6.2) is a dependency convention for the assertions which follow, not an assertion that one common threshold works for arbitrarily varying meshes.

For every positive integer $a \le M$, write uniquely

$$a = R_y(a)S_y(a), \qquad P^-(R_y(a)) > y, \qquad P^+(S_y(a)) \le y.$$

The integer $R_y(a)$, including all its prime-power multiplicities, is called the *complete rough signature* of a. A signature R is exceptional when

$$X_R := \frac{2n}{R} < X_0.$$

This complete-signature convention is important: equality of row masses then gives equality of every valuation at every prime greater than y, not merely equality of their squarefree supports.

6.1 A head-compatible balanced point

Let $K > 2$ be a fixed integer, to be chosen in a moment, and set

$$E = (2n, 2n+h], \qquad H = (2n - Kh, 2n], \qquad J = (n, 2n - Kh].$$

All endpoints are integers, so the number of integers in each half-open interval equals its Lebesgue length. Choose a small fixed $\beta > 0$, then K sufficiently large, so that

$$\alpha := \frac{h/\delta_{\mathrm{hd}} - (\beta/L)|J|}{|H|}$$

lies, for all sufficiently large n, in a fixed compact subinterval of $(0, 1/2)$. This is possible because $h = (c + o(1))n/L$: first take $\beta < c/(3\delta_{\mathrm{hd}})$, and then take a sufficiently large fixed K. By definition,

$$\delta_{\mathrm{hd}} \left(\alpha|H| + \frac{\beta}{L}|J| \right) = h \tag{6.3}$$

exactly, not only asymptotically. Define on $(n, 2n]$

$$x_a^{\mathrm{raw}} := \mathbf{1}_{(a,P_{\mathrm{hd}})=1} \left(\alpha \mathbf{1}_H(a) + \frac{\beta}{L}\mathbf{1}_J(a) \right). \tag{6.4}$$

Every flexible nonsmooth factor used below is therefore F_{hd}-free. Split $\beta = \beta_{\mathrm{prot}} + \beta_{\mathrm{act}}$ with both summands fixed and positive. On the smooth row, the two pieces of (6.4) are recorded separately as

$$\begin{aligned} x_a^{\mathrm{top}} &= \alpha \mathbf{1}_H(a)\mathbf{1}_{(a,P_{\mathrm{hd}})=1}\mathbf{1}_{P^+(a)\le y}, \\ f_a^{\mathrm{prot}} &= \frac{\beta_{\mathrm{prot}}}{L}\mathbf{1}_J(a)\mathbf{1}_{(a,P_{\mathrm{hd}})=1}\mathbf{1}_{P^+(a)\le y}, \\ z_a^{\mathrm{base}} &= \frac{\beta_{\mathrm{act}}}{L}\mathbf{1}_J(a)\mathbf{1}_{(a,P_{\mathrm{hd}})=1}\mathbf{1}_{P^+(a)\le y}. \end{aligned} \tag{6.5}$$

The first two summands will be frozen. Only the last is redistributed by the smooth bridge; (6.5) is an additive decomposition of an already normalized point, not extra mass. The de Bruijn–Saias estimate with finite head inclusion–exclusion gives

$$\sum_a z_a^{\mathrm{base}} = \{\delta_{\mathrm{hd}}\beta_{\mathrm{act}}\rho(1/\theta) + o_W(1)\}N \asymp_W N.$$

Thus this summand supplies a macroscopic active measure despite its pointwise $1/L$ scale.

There is already strong cancellation before rows are separated. If $p > W$, finite inclusion–exclusion gives, for every physical interval I and every $k \ge 1$,

$$\#\{a \in I : p^k \mid a,\ (a, P_{\mathrm{hd}}) = 1\} = \delta_{\mathrm{hd}}\frac{|I|}{p^k} + O_W(1).$$

The main terms for $\mathbf{1}_E - x^{\mathrm{raw}}$ cancel by (6.3). Hence

$$\left| \sum_a (\mathbf{1}_E(a) - x_a^{\mathrm{raw}})v_p(a) \right| \ll_W 1 + \frac{L}{\log p} = o\Big(\frac{N}{pL}\Big) \quad (W < p \le y). \tag{6.6}$$

The last estimate is uniform up to $p = y$. The work below shows that charging exceptional rows and making all other row quotas exact costs no more than the right side of (6.6) at its natural scale.

6.2 The transition-complete balanced-block estimate

We retain the de Bruijn–Saias normalization from Section 2.1, in the form proved by Hildebrand–Tenenbaum and Saias [HT93, Sai89]. For nonintegral $z > 0$, integration by parts in (2.1) gives the exact identity

$$\Lambda(z, y) = zG_y(u) - \{z\}, \qquad u = \frac{\log z}{\log y},$$

$$G_y(u) = \rho(u) - \int_0^{(u-1)_+} \rho'(u-v)\{y^v\}y^{-v}\,dv. \tag{6.7}$$

At integral z we use the right-limit convention of Theorem 2.1. For $0 < u \le 1$, the formula reduces to $\Lambda(z,y) = \lfloor z \rfloor$. Every later invocation in this section has $y = n^\theta$ and $z = tx/d$, where t ranges over a fixed compact subset of $(0,\infty)$, $d \mid P_{\mathrm{hd}}$ is fixed, and $n^{\delta_*}/C \le x \le Cn$. Hence $\log z/\log y$ remains in a fixed compact subset of $(0,\infty)$, and the compact-u error in Theorem 2.1 applies uniformly. The stated right-limit convention covers all integer physical endpoints.

Lemma 6.1 (Uniform transition normalization)**.** *On every fixed compact interval* $0 < u_- \le u \le u_+$*,*

$$|G_y(u) - G_y(v)| \ll_{u_-,u_+} |u-v|$$

uniformly for $y \ge 2$*.*

Proof. The Dickman function is Lipschitz on compact intervals. Extend ρ' by zero to the left of 1. On a fixed compact interval this extension has bounded variation (including its single jump at 1). For a compactly supported function q of bounded variation,

$$\int_{\mathbb{R}} |q(t+s) - q(t)|\,dt \le \mathrm{Var}(q)|s|.$$

Apply this to the translated copies of ρ' in (6.7), using $0 \le \{y^v\}y^{-v} \le 1$. The moving endpoint is included by the zero extension. This proves the assertion also at the transition $u = 1$. □

Lemma 6.2 (Balanced smooth blocks)**.** *Let* $x \ge n^{\delta_*}/C$*. Let a fixed number of* t_i*'s range in a fixed compact subset of* $(0,\infty)$*, and let* $a_i = O(1)$*. If*

$$\sum_i a_i t_i = 0, \qquad \sum_i |a_i| t_i |\log t_i| \ll \frac{1}{L}, \tag{6.8}$$

then

$$\sum_i a_i \Psi(t_i x, y) = O(x/L^2 + 1). \tag{6.9}$$

The same estimate holds with the smooth integers restricted to be coprime to P_{hd}*.*

Proof. Put $u = \log x/\log y$ and $d_i = \log t_i/\log y$. Theorem 2.1, (6.7), and the exact formula below y give, uniformly across every Dickman transition,

$$\Psi(t_i x, y) = t_i x G_y(u + d_i) + O\Big(1 + x\exp\{-cL^{1/2}\}\Big).$$

Subtract $t_i x G_y(u)$, sum, and use Lemma 6.1 and (6.8). This gives (6.9); the exponentially small term is absorbed because $x \ge n^{\delta_*}/C$. For the coprime version use the exact finite identity

$$\mathbf{1}_{(m,P_{\mathrm{hd}})=1} = \sum_{d \mid P_{\mathrm{hd}}} \mu(d)\mathbf{1}_{d \mid m}$$

and apply the first assertion at x/d. This proof retains every endpoint $O(1)$, and differentiates neither ρ nor Ψ at a transition. □

For later use isolate the effect of a fixed head divisor on the entire physical block. Put

$$B_{\mathrm{ph}}(x) := \alpha\{\Psi(x,y) - \Psi((1-K\kappa)x,y)\} + \frac{\beta}{L}\{\Psi((1-K\kappa)x,y) - \Psi(x/2,y)\}.$$

Lemma 6.3 (Fixed-head shift of a physical block)**.** *Uniformly for* $n^{\delta_*}/C \le x \le Cn$ *and every fixed* $d \mid P_{\mathrm{hd}}$*,*

$$B_{\mathrm{ph}}(x/d) = \frac{1}{d}B_{\mathrm{ph}}(x) + O_W(x/L^2 + 1). \tag{6.10}$$

Proof. Write $b = 1 - K\kappa$, $u = \log x/\log y$, $e = \log d/\log y$, and

$$\Delta_e(s) := G_y(s-e) - G_y(s).$$

The fixedness of d, Lemma 6.1, and the triangle inequality give, on the relevant compact parameter interval,

$$|\Delta_e(s)| \ll_W L^{-1}, \qquad |\Delta_e(s) - \Delta_e(t)| \ll_W |s-t|. \tag{6.11}$$

The transition normalization used in the proof of Lemma 6.2, with every endpoint floor retained, therefore gives

$$\begin{aligned} B_{\rm ph}(x/d) - d^{-1}B_{\rm ph}(x) = \frac{x}{d}\Big[&\alpha\{\Delta_e(u) - b\Delta_e(u+\log b/\log y)\} \\ &+\frac{\beta}{L}\{b\Delta_e(u+\log b/\log y) - \tfrac12\Delta_e(u-\log 2/\log y)\}\Big] \\ &+ O_W(xe^{-cL^{1/2}}+1). \end{aligned} \tag{6.12}$$

Here $1-b = K\kappa = O(L^{-1})$ and $|\log b|/\log y = O(L^{-2})$. Rewrite the first brace as

$$(1-b)\Delta_e(u) + b\{\Delta_e(u) - \Delta_e(u+\log b/\log y)\}.$$

Both of its terms are $O_W(L^{-2})$ by (6.11); the second brace in (6.12) is $O_W(L^{-1})$ and already has the prefactor L^{-1}. The exponential error is $O(x/L^2)$ in the stated range. This proves (6.10), including the $O(1)$ endpoint terms. $\square$

Fix a nonexceptional signature R, put $x = X_R$, and write $\kappa = h/(2n)$. The exact upper-row count is

$$t_R = \Psi((1+\kappa)x, y) - \Psi(x,y).$$

The mass of (6.4) in this row is the corresponding $P_{\rm hd}$-coprime smooth count in $((1-K\kappa)x, x]$ with coefficient α, plus that in $(x/2, (1-K\kappa)x]$ with coefficient β/L. Dividing (6.3) by R gives the exact real-length identity

$$\delta_{\rm hd}\left\{\alpha K\kappa x + \frac{\beta}{L}(1/2 - K\kappa)x\right\} = \kappa x. \tag{6.13}$$

The head inclusion–exclusion is handled before invoking the balanced lemma. Exactly from the definition of $B_{\rm ph}$,

$$m_R^{\rm raw} = \sum_{d|P_{\rm hd}} \mu(d) B_{\rm ph}(x/d).$$

By Lemma 6.3,

$$m_R^{\rm raw} = \delta_{\rm hd}B_{\rm ph}(x) + O_W(x/L^2+1). \tag{6.14}$$

It remains to compare the unscaled upper block with $\delta_{\rm hd}B_{\rm ph}(x)$. Its endpoint pairs are

$$\Psi((1+\kappa)x,y) - \Psi(x,y) - \delta_{\rm hd}B_{\rm ph}(x).$$

Their coefficients a_i and relative endpoints t_i have $\sum_i a_i t_i = 0$ by (6.13). Moreover, the endpoints $1+\kappa$ and $1-K\kappa$ contribute $O(\kappa) = O(L^{-1})$ to $\sum_i |a_i||t_i|\,|\log t_i|$, the endpoint 1 contributes zero, and the fixed endpoint $1/2$ has coefficient $\delta_{\rm hd}\beta/L$. Thus the absolute-log hypothesis (6.8) really does hold for this unscaled comparison. Applying Lemma 6.2 and then (6.14), with $m_R^{\rm raw} = \sum_{R_y(a)=R} x_a^{\rm raw}$,

$$\boxed{\varepsilon_R := t_R - m_R^{\rm raw} = O_W(X_R/L^2+1)} \qquad (X_R \ge X_0). \tag{6.15}$$

The $+1$ in (6.15) includes the integer endpoint error of every one of the finitely many head inclusion–exclusion terms. The same argument over the fixed-ratio interval J/R, together with the positivity of $G_y(u) = \rho(u) + O(1/L)$ on the relevant compact range, gives

$$\#\{a \in J : R_y(a) = R,\ (a, P_{\rm hd}) = 1\} \gg_W X_R. \tag{6.16}$$

6.3 The fixed-depth exceptional expansion

The estimate needed for exceptional rows lies strictly between the fourth and fifth births of a rough integer. We give the short proof because it is what permits the signed sum below; a rowwise sawtooth estimate would not be uniform at the one-prime boundary.

For an interval I, write

$$\Phi(I;y) := \#\{r \in I : P^-(r) > y\}.$$

Lemma 6.4 (Four-to-five rough chamber)**.** *For each fixed $C \geq 1$, let Z satisfy*

$$4.2 \leq \frac{\log Z}{\log y} \leq 4.6,$$

and let $I \subset [Z/C, CZ] \cap (0, 3n]$ have arbitrary real endpoints. There is a C^1 function $\mathcal{K}$ on $[4.1, 4.7]$, with bounded derivative, such that

$$\Phi(I;y) = \int_I \frac{\mathcal{K}(\log t/\log y)}{\log y}\,dt + O_C\Big(Ze^{-c\sqrt{L}} + \sqrt{Z}\Big). \tag{6.17}$$

Proof. Changing the inclusion convention at either endpoint of I changes $\Phi(I;y)$ by at most two, which is absorbed by the claimed error. We may therefore write $I = (A, B]$. For $1 \leq j \leq 4$, define the ordered count

$$\mathcal{N}_j(I;y) := \#\{(p_1, \ldots, p_j) : p_i > y \text{ prime},\ p_1 \cdots p_j \in I\}.$$

We first record the reciprocal-prime discrepancy with all constants uniform on the compact range used below. On $[1, 4.8]$, put

$$d\mu_y(s) := \sum_{p>y} \frac{1}{p}\delta_{\log p/\log y}(ds), \qquad d\lambda(s) := \frac{ds}{s}.$$

For $1 \leq T \leq 4.8$, partial summation gives

$$\mu_y([1,T]) = \frac{\pi(y^T)}{y^T} - \frac{\pi(y)}{y} + \int_y^{y^T} \frac{\pi(v)}{v^2}\,dv. \tag{6.18}$$

Write $E(v) = \pi(v) - \mathrm{li}(v)$. The li-part of (6.18) is exactly $\int_y^{y^T} dv/(v\log v) = \log T$. By (2.7), uniformly in T,

$$\begin{aligned} &\left|\frac{E(y^T)}{y^T}\right| + \left|\frac{E(y)}{y}\right| + \int_y^{y^{4.8}} \frac{|E(v)|}{v^2}\,dv \\ &\quad \ll e^{-c_0\sqrt{\log y}} + \int_y^{y^{4.8}} e^{-c_0\sqrt{\log v}}\frac{dv}{v} \ll Le^{-c_0\sqrt{(2/9)L}} \ll e^{-c_1\sqrt{L}}. \end{aligned}$$

Consequently

$$D_y := \sup_{1\leq T\leq 4.8} |\mu_y([1,T]) - \lambda([1,T])| \ll e^{-c_1\sqrt{L}}.$$

The same bound holds for either convention at the moving endpoint: changing one atom costs at most $1/y \ll e^{-c_1\sqrt{L}}$. In particular, if $1 \leq a \leq 4.8$ and f has bounded variation on $[1, a]$, Stieltjes integration by parts gives

$$\left|\int_{[1,a]} f\,d(\mu_y - \lambda)\right| \leq 2(D_y + y^{-1})\{\|f\|_\infty + \mathrm{Var}_{[1,a]}\, f\}. \tag{6.19}$$

We next expose every moving endpoint in the ordered count. Fix j, put $m = j-1$, and write $q = p_1 \cdots p_m$, with the empty product interpreted as one. If $B/q \le y$, this tuple contributes nothing. Otherwise put $a_q = \max(y, A/q)$. The two nonempty endpoints $a_q, B/q$ lie between $y = n^{2/9}$ and $CZ \le Cy^{4.6}$. Subtracting (2.7) at these exact endpoints gives

$$\begin{aligned}\#\{p_j > y : qp_j \in (A,B]\} &= \pi(B/q) - \pi(a_q) \\ &= \int_{a_q}^{B/q} \frac{dv}{\log v} + O_C\Big(\frac{Z}{q} e^{-c_2\sqrt{L}}\Big) \\ &= \int_A^B \frac{\mathbf{1}_{qy<t}}{q\log(t/q)}\, dt + O_C\Big(\frac{Z}{q} e^{-c_2\sqrt{L}}\Big). \end{aligned} \tag{6.20}$$

This is an absolute cumulative-PNT error, not a relative assertion for a short interval. Moreover,

$$\sum_{\substack{p_1,\dots,p_m>y \\ q<B/y}} \frac{1}{q} \le \Bigg(\sum_{y<p\le CZ/y} \frac{1}{p}\Bigg)^m \ll_{C,m} 1$$

by (2.10). Thus the total error in (6.20) is $O_{C,j}(Ze^{-c_2\sqrt{L}})$, uniformly in both endpoints.

Summing the main term in (6.20) and interchanging the finite sum with the t-integral gives

$$\mathcal{N}_j(I;y) = \int_I S_{j,y}(t)\, dt + O_{C,j}(Ze^{-c_2\sqrt{L}}). \tag{6.21}$$

If $u = \log t / \log y$, then

$$S_{j,y}(t) = \frac{1}{\log y} \int_{\substack{s_1,\dots,s_m \ge 1 \\ s_1+\cdots+s_m<u-1}} \frac{d\mu_y(s_1)\cdots d\mu_y(s_m)}{u - s_1 - \cdots - s_m}; \tag{6.22}$$

for $m=0$, the last integral means $1/u$. The strict face causes no outer-integral error: for each fixed prime tuple it is a single value of t.

For $t \in [Z/C, CZ]$, the hypotheses on Z imply, for all sufficiently large n depending only on C,

$$4.1 \le u \le 4.7.$$

Replace the $m \le 3$ copies of $d\mu_y$ in (6.22) one at a time by $d\lambda$. Conditioned on the other variables, the current variable ranges over $[1,a]$, where $a = u - 1 - \sum_{\ell \ne i} s_\ell$, or over the empty set. On a nonempty interval the conditional kernel

$$f(s) = \frac{1}{u - s - \sum_{\ell\ne i} s_\ell}$$

has denominator at least one, and hence

$$\|f\|_\infty \le 1, \qquad \operatorname{Var} f \le 1.$$

For large n, both μ_y and λ have mass at most $M_0 := 1 + \log 4.8$ on the relevant compact interval. Therefore the literal m-term telescope and (6.19) give

$$\begin{aligned}\Bigg| \int_{\sum s_i < u-1} \frac{d\mu_y(s_1)\cdots d\mu_y(s_m)}{u - \sum s_i} - \int_{\sum s_i \le u-1} \frac{ds_1\cdots ds_m}{s_1\cdots s_m(u-\sum s_i)} \Bigg| \\ \le 4mM_0^{m-1}(D_y + y^{-1}) \ll_j e^{-c_1\sqrt{L}}. \end{aligned} \tag{6.23}$$

This includes a moving or degenerate simplex face with one constant.

Define

$$\mathcal{K}_1(u) = \frac{1}{u},$$
$$\mathcal{K}_j(u) = \int_{\substack{s_1,\dots,s_{j-1}\ge 1\\ s_1+\cdots+s_{j-1}\le u-1}} \frac{ds_1\cdots ds_{j-1}}{s_1\cdots s_{j-1}(u-s_1-\cdots-s_{j-1})} \qquad (2\le j\le 4).$$

Equations (6.21) and (6.23) prove

$$\mathcal{N}_j(I;y) = \frac{1}{\log y}\int_I \mathcal{K}_j(\log t/\log y)\,dt + O_{C,j}(Ze^{-c\sqrt{L}}). \tag{6.24}$$

For completeness, the regularity of the kernels is also uniform. For $j\ge 2$, put $m=j-1$, $T=u-j$, and $\Delta_m=\{z_i\ge 0:\sum z_i\le 1\}$. The substitutions $s_i=1+Tz_i$ give

$$\mathcal{K}_j(u) = T^m\int_{\Delta_m}\frac{dz_1\cdots dz_m}{\prod_{i=1}^m(1+Tz_i)\{1+T(1-\sum_i z_i)\}}.$$

On $4.1\le u\le 4.7$ and $2\le j\le 4$, one has $T\ge 0.1$ and $T\ll 1$. The integrand and its first T-derivative are bounded on the fixed simplex, so differentiation under the integral proves

$$\sup_{4.1\le u\le 4.7}\{|\mathcal{K}_j(u)|+|\mathcal{K}_j'(u)|\}\ll_j 1.$$

The same assertion is immediate for $\mathcal{K}_1$.

It remains to pass from ordered tuples to rough integers. Put $\Omega(r)=\sum_p v_p(r)$, and let $A_j(I;y)$ count the $r\in I$ with $P^-(r)>y$ and $\Omega(r)=j$. Since $y^5=n^{10/9}>3n$, every integer counted by $\Phi(I;y)$ has $1\le\Omega(r)\le 4$. If

$$r=\prod_{\nu=1}^k q_\nu^{e_\nu}, \qquad \sum_{\nu=1}^k e_\nu = j,$$

then its exact multiplicity in $\mathcal{N}_j$ is

$$\frac{j!}{e_1!\cdots e_k!}.$$

Thus division by $j!$ gives weight one for a squarefree r, and weight $1/(e_1!\cdots e_k!)\in(0,1]$ otherwise. Consequently

$$0\le\Phi(I;y)-\sum_{j=1}^4\frac{\mathcal{N}_j(I;y)}{j!}\le\#\{r\in I: q^2\mid r \text{ for some prime } q>y\}. \tag{6.25}$$

The last quantity is at most

$$\sum_{\substack{y<q\le\sqrt{CZ}\\ q\ \text{prime}}}\left(\frac{|I|}{q^2}+1\right)\ll_C\frac{|I|}{y\log y}+\frac{\sqrt{Z}}{\log Z}$$
$$\ll_C Ze^{-c\sqrt{L}}+\sqrt{Z}.$$

Here the first bound follows by partial summation from $\pi(t)\ll t/\log t$, and the last uses $y^{-1}\ll e^{-c\sqrt{L}}$. Combining (6.24) and (6.25) proves (6.17) with

$$\mathcal{K}(u)=\sum_{j=1}^4\frac{\mathcal{K}_j(u)}{j!}.$$

□

For later use we record explicitly the simultaneous form just proved. For fixed J_0, G, let $J \le J_0$, let $I_\nu \subset [Z/C, CZ] \cap (0, 3n]$ be intervals, possibly empty or clipped at arbitrary moving endpoints, and let $|\gamma_\nu| \le G$. Then the same kernel and one common implied constant give

$$\sum_{\nu=1}^{J} \gamma_\nu \Phi(I_\nu; y) = \frac{1}{\log y} \sum_{\nu=1}^{J} \gamma_\nu \int_{I_\nu} \mathcal{K}(\log t/\log y)\, dt + O_{C,J_0,G}(Ze^{-c\sqrt{L}} + \sqrt{Z}). \tag{6.26}$$

Put $R_0 = 2n/X_0$, and let

$$w_{\mathrm{hd}}(a) := \mathbf{1}_E(a) - x_a^{\mathrm{raw}}.$$

For $W < p \le y$, the exact signed contribution of the exceptional rows to the uncharged comparison is

$$\mathcal{E}_p = \sum_{a: R_y(a) > R_0} w_{\mathrm{hd}}(a) v_p(a) = \sum_{b < 2X_0} v_p(b) \sum_{\substack{P^-(R) > y \\ R > R_0}} w_{\mathrm{hd}}(bR). \tag{6.27}$$

Indeed, $a = Rb$ in this range has $b < 2X_0 < y$, and so $v_p(a) = v_p(b)$. Moreover, throughout the support in the inner sum,

$$4.25 + o(1) < \frac{\log R}{\log y} < 4.51. \tag{6.28}$$

The lower bound follows from $(1 - \delta_*)/\theta > 17/4$, and the upper bound from $R < 3n$. Thus Lemma 6.4 applies with a fixed separation from both birth faces.

Define the mean-zero periodic function

$$c_{\mathrm{hd}}(b) := 1 - \delta_{\mathrm{hd}}^{-1} \mathbf{1}_{(b, P_{\mathrm{hd}}) = 1}.$$

For every $b < 2X_0$, put

$$Z_b := \frac{2n}{b}, \qquad u_b := \frac{\log Z_b}{\log y}, \qquad g_b := \mathbf{1}_{(b, P_{\mathrm{hd}}) = 1},$$

and introduce the three disjoint physical intervals

$$\begin{aligned} I_{+,b} &= (Z_b, (1+\kappa) Z_b], \\ I_{-,b} &= ((1 - K\kappa) Z_b, Z_b], \\ I_{0,b} &= (Z_b/2, (1 - K\kappa) Z_b]. \end{aligned}$$

Since $b < 2X_0 < y$, the complete rough factor is coprime to P_{hd}. The definitions of E, H, J and (6.4) therefore give the exact identity

$$\sum_{P^-(R) > y} w_{\mathrm{hd}}(bR) = \Phi(I_{+,b}; y) - g_b \alpha \Phi(I_{-,b}; y) - g_b \frac{\beta}{L} \Phi(I_{0,b}; y). \tag{6.29}$$

Suppose first that $b \le X_0/2$. The cutoff $R > R_0$ in (6.27) is then automatic. All three intervals lie in $[Z_b/2, 2Z_b] \cap (0, 3n]$, and the range in (6.28) puts u_b in the compact chamber of Lemma 6.4. With

$$\mathfrak{E}_b := Z_b e^{-c\sqrt{L}} + \sqrt{Z_b},$$

the simultaneous estimate (6.26), applied once to the three terms of (6.29), yields

$$\begin{aligned} \sum_{P^-(R) > y} w_{\mathrm{hd}}(bR) = \frac{1}{\log y} \Bigg\{ &\int_{I_{+,b}} \mathcal{K}\Big(\frac{\log t}{\log y}\Big)\, dt - g_b \alpha \int_{I_{-,b}} \mathcal{K}\Big(\frac{\log t}{\log y}\Big)\, dt \\ &- g_b \frac{\beta}{L} \int_{I_{0,b}} \mathcal{K}\Big(\frac{\log t}{\log y}\Big)\, dt \Bigg\} + O_W(\mathfrak{E}_b). \end{aligned} \tag{6.30}$$

This is one endpoint-uniform invocation, rather than three asymptotics with potentially different moving-endpoint constants.

Put $K_b = \mathcal{K}(u_b)$. By the mean-value theorem and the uniform bound for $\mathcal{K}'$, the error made by replacing every kernel in (6.30) by K_b is at most

$$\frac{1}{\log y}\left\{\int_{I_{+,b}}\left|\mathcal{K}\left(\frac{\log t}{\log y}\right)-K_b\right|dt+\alpha\int_{I_{-,b}}\left|\mathcal{K}\left(\frac{\log t}{\log y}\right)-K_b\right|dt\right.$$
$$\left.+\frac{\beta}{L}\int_{I_{0,b}}\left|\mathcal{K}\left(\frac{\log t}{\log y}\right)-K_b\right|dt\right\}\ll_W \frac{Z_b}{(\log y)^2}\left(\kappa^2+K^2\kappa^2+\frac{1}{L}\right)\ll_W \frac{Z_b}{L^3}.$$

Indeed, the two short intervals have respectively relative lengths κ and $K\kappa$, while $|\log(t/Z_b)|\leq \log 2$ on the broad interval; also $\kappa \ll L^{-1}$.

The constant-kernel coefficient has no asymptotic error. Dividing (6.3) by b, or equivalently using (6.13) with $x = Z_b$, gives

$$\kappa Z_b - g_b\left\{\alpha K\kappa Z_b+\frac{\beta}{L}(1/2-K\kappa)Z_b\right\}$$
$$=\kappa Z_b\left(1-\frac{g_b}{\delta_{\rm hd}}\right)=\frac{h}{b}\,c_{\rm hd}(b).$$

Finally, $Z_b \geq 4n/X_0$ in the present range, so

$$\mathfrak{E}_b \ll Z_b/L^3. \tag{6.31}$$

Combining (6.30)– (6.31), and retaining a harmless unit endpoint allowance, proves

$$\sum_{P^-(R)>y} w_{\rm hd}(bR)=\frac{h}{b\log y}\mathcal{K}\left(\frac{\log(2n/b)}{\log y}\right)c_{\rm hd}(b)+O_W\left(\frac{n}{bL^3}+1\right).$$

For $p > W$, multiplication by p^k permutes the residue classes modulo $P_{\rm hd}$. Hence

$$\sup_X\left|\sum_{m\leq X} c_{\rm hd}(p^k m)\right|\ll_W 1.$$

Partial summation, followed by the expansion $v_p(b)=\sum_{k\geq 1}\mathbf{1}_{p^k|b}$, shows that for every bounded C^1 function V of $\log b/\log y$,

$$\sum_{b\leq X}\frac{c_{\rm hd}(b)v_p(b)}{b}V\left(\frac{\log b}{\log y}\right)\ll_W \frac{1}{p}. \tag{6.32}$$

For the single cutoff band $X_0/2 < b < 2X_0$, set

$$I^\circ_{\nu,b} := I_{\nu,b}\cap(R_0,\infty) \qquad (\nu\in\{+,-,0\}).$$

The three physical intervals are disjoint, so taking the absolute value of each coefficient is exact:

$$\sum_{\substack{P^-(R)>y\\ R>R_0}} |w_{\rm hd}(bR)|$$
$$=\Phi(I^\circ_{+,b};y)+g_b\alpha\Phi(I^\circ_{-,b};y)+g_b\frac{\beta}{L}\Phi(I^\circ_{0,b};y).$$

Every clipped interval is still contained in $[Z_b/2, 2Z_b]\cap(0,3n]$. Applying (6.26) simultaneously, now with positive coefficients, gives its main term bounded by

$$\frac{\|\mathcal{K}\|_\infty}{\log y}\left\{|I^\circ_{+,b}|+\alpha|I^\circ_{-,b}|+\frac{\beta}{L}|I^\circ_{0,b}|\right\}\leq\frac{\|\mathcal{K}\|_\infty}{\log y}\left\{\kappa Z_b+\alpha K\kappa Z_b+\frac{\beta Z_b}{2L}\right\}$$

$$\ll_W \frac{Z_b}{L^2}.$$

Uniformly in this band, $Z_b \asymp n/X_0$, and hence the common analytic error $Z_b e^{-c\sqrt{L}} + \sqrt{Z_b}$ is $o(Z_b/L^2)$. Since $Z_b = 2n/b$, we obtain, with the existing harmless endpoint allowance,

$$\sum_{\substack{P^-(R)>y \\ R>R_0}} |w_{\mathrm{hd}}(bR)| \ll_W \frac{n}{bL^2} + 1.$$

Finally, the elementary prime-power identities

$$\sum_{b\le X} \frac{v_p(b)}{b} \ll \frac{\log X + 1}{p}, \qquad \sum_{b\le X} v_p(b) \ll \frac{X}{p},$$
$$\sum_{X/2<b<2X} \frac{v_p(b)}{b} \ll \frac{1}{p}$$

make the last summation completely explicit. The function $b \mapsto \mathcal{K}(\log(2n/b)/\log y)$ is a bounded C^1 function of $\log b/\log y$. Hence (6.32) and $h \ll n/L$ give

$$\frac{h}{\log y}\left|\sum_{b\le X_0/2} \frac{c_{\mathrm{hd}}(b)v_p(b)}{b}\mathcal{K}\left(\frac{\log(2n/b)}{\log y}\right)\right| \ll_W \frac{n}{pL^2}.$$

The accumulated error in the deep range is, with no suppressed logarithm,

$$\sum_{b\le X_0/2} v_p(b)\left(\frac{n}{bL^3} + 1\right) \ll \frac{n}{L^3}\frac{\log X_0 + 1}{p} + \frac{X_0}{p} \ll \frac{n}{pL^2} + \frac{X_0}{p}.$$

Likewise, the entire cutoff band contributes

$$\sum_{X_0/2<b<2X_0} v_p(b)\left(\frac{n}{bL^2} + 1\right) \ll \frac{n}{pL^2} + \frac{X_0}{p}.$$

Finally,

$$\frac{X_0}{p} = o\left(\frac{n}{pL^2}\right)$$

uniformly in p, because $X_0 = n^{\delta_*}$ and $\delta_* < 1$. Applying these three ledgers to (6.27), we obtain

$$\boxed{|\mathcal{E}_p| \ll_W \frac{n}{pL^2} = \frac{N}{pL}} \qquad (W < p \le y). \tag{6.33}$$

The order of summation in (6.27) is essential: absolute values are taken only after the balanced physical pieces have been summed for a fixed smooth core b. Thus every one-sided rough birth remains in the displayed signed sum.

6.4 Exact row correction, including all endpoint terms

We first describe the correction without guards. In each nonexceptional, nontrivial row choose a clean subpool $\mathcal{P}_R \subset J$ as in (6.16), and let μ_R be its uniform probability measure. Adding

$$\nu_R(a) = \varepsilon_R \mu_R(a) \tag{6.34}$$

makes the row mass exactly t_R. Equations (6.15) and (6.16) give

$$\|\nu_R\|_\infty \ll_W L^{-2} + X_R^{-1} = o(L^{-1}), \tag{6.35}$$

uniformly because $X_R \geq X_0$. Thus the broad floor remains between fixed positive multiples of $1/L$, and every weight stays strictly between zero and one.

For $W < p \leq y$, the number of core candidates in $\mathcal{P}_R$ divisible by p^k is at most $O_W(X_R/p^k+1)$. If it is nonzero, then $p^k \ll_W X_R$, so the endpoint $+1$ is itself $O_W(X_R/p^k)$. Consequently

$$\sum_a \mu_R(a)v_p(a) \ll_W \frac{1}{p}.$$

The rough harmonic and counting bounds

$$\sum_{\substack{R\leq 2n/X_0\\ P^-(R)>y}} \frac{1}{R} \ll_\theta 1, \qquad \#\{R \leq 2n/X_0 : P^-(R) > y\} \ll_\theta \frac{N}{X_0}$$

follow respectively from the truncated Euler product and the elementary upper-bound sieve. Therefore

$$\begin{aligned}\sum_R \left|\sum_a \nu_R(a)v_p(a)\right| &\ll_W \frac{1}{p}\sum_R (X_R/L^2+1)\\ &\ll_W \frac{1}{p}\left(\frac{n}{L^2}+\frac{N}{X_0}\right) \ll_W \frac{N}{pL}.\end{aligned} \tag{6.36}$$

This computation includes all prime powers and all $+1$'s. For an independent check, group rows dyadically by $X_R \asymp X$. There are $O_\theta(N/X)$ such signatures. If an endpoint $+1$ is retained instead of being absorbed into X_R/p^k, only scales $X \gg \max(X_0, p)$ occur, and its sum is

$$\ll \frac{L}{\log p} \sum_{X \gg \max(X_0,p)} \frac{N}{X}\left(\frac{1}{L^2}+\frac{1}{X}\right) \ll \frac{N}{pL};$$

here either X_0 or p dominates every power of L.

We now insert the charged guard system. Define the exceptional upper list and the global bank-donor list by

$$E_{\text{exc}} := \{a \in E : X_{R_y(a)} < X_0\}, \qquad E_{\text{donor}} := \{a \in E : a \text{ is a designated bank donor}\}.$$

Let G^0_{bank} be the chosen state-zero bank factors and define, exhaustively,

$$G_{\text{fix}} := E_{\text{exc}} \setminus E_{\text{donor}}. \tag{6.37}$$

There is no further class of dedicated fixed residual tokens: anchors are external to the residual product, bank states belong to G^0_{bank}, and all later bridge and tangent factors remain flexible. Put

$$B_0 := \prod_{a\in G_{\text{fix}}} a \prod_{a\in G^0_{\text{bank}}} a. \tag{6.38}$$

Also write $G^1_{\text{bank}} := \bigcup_{g\in\mathcal{B}} G^1_g$. The following ownership ledger distinguishes selection classes from source or guard roles; this matters because a bottom donor can be the same occurrence as one of its two bank states. The symbol $\mathcal{A} = \coprod_R \mathcal{A}_R$ will denote the lower coordinates left flexible after these exclusions. To avoid conflict with the quotient space denoted by $\mathcal{G}$ in Section 8, denote the exhaustive set of numerical guards by

$$\Gamma_{\text{num}} := \mathcal{H}' \cup G_{\text{fix}} \cup E_{\text{donor}} \cup G^0_{\text{bank}} \cup G^1_{\text{bank}}. \tag{6.39}$$

Object	Location	Selection status	Row-quota status	Later operations
$\mathcal{H}'$	$(n, 2n]$	fixed external anchors	outside the residual quotas	never varied
E_{exc}	$(2n, M]$	source list, not itself a selection class	supplies exceptional upper tokens	split into G_{fix} and donor occurrences
E_{donor}	$(2n, M]$	withheld backing token; not inserted separately	one token is replaced by its bank state	never bridged or rounded
G_{fix}	$(2n, M]$	fixed in B_0 and in the final residual set	charged before q_R is frozen	never varied
$G^0_{\mathrm{bank}}, G^1_{\mathrm{bank}}$	$(n, M]$	alternative full path states; state 0 is in B_0	equal complete-signature counts	toggled only after rounding
$\mathcal{A}$	$(n, 2n]$	flexible candidates	supplies the remaining quota q_R	bridge, tangent, and rounding
Γ_{num}	$(n, M]$	exclusion role, not an additional selection class	no additional quota beyond its constituent roles	used only to enforce disjointness

The final selection classes are $\mathcal{H}'$, G_{fix}, one full state of each bank path, and the rounded subset of $\mathcal{A}$. They are pairwise disjoint. The only allowed overlap in the table is a role overlap: in the last two bottom pools an element of E_{donor} may be a bank-state occurrence, and it is then counted once, as stipulated in Lemma 5.1.

Lemma 6.5 (Guard census)**.** *Before the quotas q_R are frozen, the numerical exclusions are exhausted by the external anchors, the exceptional source and donor decisions, and both states of the reserved bank. They have the following two, logically different, size properties.*

1. *Globally,*

$$|\mathcal{H}'| = O_c(N), \quad |G_{\mathrm{fix}}| = O_c(N), \quad |E_{\mathrm{donor}}| + |G^0_{\mathrm{bank}}| + |G^1_{\mathrm{bank}}| = O(yL).$$

   Only $O(y)$ anchors are modified, and the number of promoted anchors having trivial rough signature is $O(\pi(y))$.

2. *In any active nonexceptional row $R \neq 1$, at most one anchor and at most one bank component occur; G_{fix} contributes nothing. Hence deletion of all guarded lower coordinates removes $O(1)$ candidates from that row. For a fixed medium-prime pair u, v, the only anchors which can equal ua or va are the $O(\pi(y))$ promoted smooth anchors; together with the bank endpoints, they give $O(yL)$ potentially relevant lower guards.*

In particular, the proof uses a local rowwise census and a common-list-relevant census. It does not require the stronger statement that all anchors and all fixed exceptional factors together number $O(yL^A)$, which need not hold.

Proof. The first global bounds follow from the construction of the central anchors, $|E| = h = O(N)$, and (5.14). Anchor modification affects $O(y)$ occurrences by the fixed-prefix scale count (5.15). A promoted anchor has trivial rough signature only when its base prime is at most y, which gives $O(\pi(y))$ possibilities.

A nonsmooth prefix or row-zero anchor has one unique rough marker $P > y$; a promoted nonsmooth anchor has one complete signature $p^{e_p} > 1$. There is at most one such anchor in a given nontrivial row. Distinct bank components have distinct marker primes by Lemma 5.1, so a nontrivial marker row contains at most one component, its two possible endpoint occurrences, and its one donor occurrence. The set G_{fix} is supported only on exceptional rows. This proves the rowwise assertion.

Finally, a prefix anchor has smooth cofactor supported on $F_{\mathrm{anc}} \subset [2, W]$, a row-zero anchor is prime, and a promoted nonsmooth anchor has a power of two as smooth cofactor. None is divisible by a label in $(W, y]$. Only promoted smooth anchors remain possible in a lower common-multiplier list. The fixed exceptional factors and all withheld donors lie above $2n$; bank endpoints contribute $O(yL)$ numbers. This proves the last assertion. □

All factors in this product are mutually distinct and are guarded from the anchors and from the flexible candidates. If $m_R(G)$ denotes the number of factors of G having complete signature R, define, for every nontrivial row,

$$q_R := t_R - m_R(G_{\mathrm{fix}}) - m_R(G^0_{\mathrm{bank}}) \in \mathbb{Z}_{\geq 0}. \tag{6.40}$$

The backing map is literal:

$$\iota(g) = g \quad (g \in G_{\mathrm{fix}}), \qquad \iota(g) = \text{the designated donor of its bank component} \quad (g \in G^0_{\mathrm{bank}}). \tag{6.41}$$

Its image lies in E, preserves the complete signature, and is injective. Indeed, different bank components have different donors, and every such donor has been removed from G_{fix} (a nonexceptional donor never belonged to G_{fix} in the first place). This proves the nonnegativity of (6.40). In an exceptional row, the partition into retained factors and designated donors, together with the one base token replacing each donor, exhausts all t_R upper tokens; hence $q_R = 0$ there.

We also make the whole-row convention explicit. Let $\mathscr{B}_{\mathrm{row}}$ be the set of nontrivial complete signatures of bank component marker rows. Every such signature is the component's unique marker prime $P > y$. Its designated donor is $Pu \in E$, with $u \ll y$, so

$$X_P = \frac{2n}{P} \ll y.$$

The component count (5.14) and the use of a different marker for every component give

$$|\mathscr{B}_{\mathrm{row}}| = O(yL), \qquad X_R \ll y \quad (R \in \mathscr{B}_{\mathrm{row}}).$$

In the endpoint-deletion convention used in this paper, every nonexceptional bank-touched row remains active and is corrected below; no such row is dedicated in its entirety. Thus the set of fully dedicated nonexceptional rows is

$$\mathscr{D}_{\mathrm{row}} := \varnothing \subset \mathscr{B}_{\mathrm{row}}, \qquad |\mathscr{D}_{\mathrm{row}}| = O(yL), \qquad X_R \ll y \quad (R \in \mathscr{D}_{\mathrm{row}}). \tag{6.42}$$

The last two bounds would remain true if any subset of the bank-touched rows were instead dedicated. Exceptional rows $X_R < X_0$ are fixed separately by (6.37) and are not put in $\mathscr{D}_{\mathrm{row}}$.

Lemma 6.6 (Uniform clean pool after all guards)**.** *For every active nonexceptional row $R \neq 1$, let $\mathcal{A}_R$ be the lower candidates remaining after the exhaustive guard census in Lemma 6.5, and let $\mathcal{P}_R \subset \mathcal{A}_R$ be the broad correction pool. Uniformly in these rows,*

$$\Bigg| q_R - \sum_{a \in \mathcal{A}_R} x_a^{\mathrm{raw}} \Bigg| \ll_W X_R/L^2 + 1, \qquad \#\mathcal{P}_R \gg_W X_R. \tag{6.43}$$

Consequently the constant correction on $\mathcal{P}_R$ has density

$$O_W(L^{-2} + X_R^{-1}) = o(L^{-1}), \tag{6.44}$$

and the two-sided raw slack survives.

Proof. Before guards, (6.16) supplies at least $c_W X_R$ broad candidates. By Lemma 6.5, at most a fixed number are deleted in an active nontrivial row: at most one anchor, and the endpoints and donor role of at most one bank component. Fixed exceptional factors lie above $2n$ and occur only in exceptional rows. Since $X_R \geq X_0 \to \infty$, at least $c_W X_R/2$ clean candidates remain for all sufficiently large n, proving the second assertion in (6.43).

For the first assertion, let $\mathcal{D}_R$ be exactly the lower coordinates removed from the raw row. Subtracting the definitions of q_R and $\mathcal{A}_R$ gives the identity

$$\left(q_R - \sum_{a\in\mathcal{A}_R} x_a^{\mathrm{raw}}\right) - \left(t_R - \sum_{R_y(a)=R} x_a^{\mathrm{raw}}\right)$$
$$= -m_R(G_{\mathrm{fix}}) - m_R(G^0_{\mathrm{bank}}) + \sum_{a\in\mathcal{D}_R} x_a^{\mathrm{raw}}.$$

Here $m_R(G_{\mathrm{fix}}) = 0$. If a bank component occurs in the row, its contribution to the right side is

$$-1 + \sum_{\substack{a\in\mathcal{D}_R \\ a \text{ is one of its lower endpoints}}} x_a^{\mathrm{raw}} = O_W(1);$$

the expression is not claimed to vanish. The at-most-one anchor contributes another $O_W(1)$. There are no further classes in the census. Therefore

$$\left| q_R - \sum_{a\in\mathcal{A}_R} x_a^{\mathrm{raw}} \right| \leq |\varepsilon_R| + O_W(1),$$

and (6.15) proves the first assertion. Dividing by $\#\mathcal{P}_R \gg_W X_R$ gives (6.44). Since $X_R \geq n^{\delta_*}$, this is $o(L^{-1})$, so a fixed fraction of both the lower and upper broad margin remains. $\square$

The already charged medium-prime guard profile is $O_W(N/(pL))$. Replacing ε_R in (6.34) by the left side of (6.43) therefore repeats (6.35)–(6.36) verbatim and proves the exact postcharge identity

$$\sum_{a\in\mathcal{A}_R} x_a = q_R \qquad (R \neq 1). \tag{6.45}$$

The row $R = 1$ is intentionally left for the integral smooth-row quota and structured head fit of Section 8; its top component and protected floor in (6.5) remain frozen.

6.5 Nonnegative exceptional charge and simultaneous feasibility

Remark 6.7 (Signed discrepancy is not a charge). The quantity $\mathcal{E}_p$ in (6.27) is signed: it compares the actual upper exceptional rows with their lower raw baseline and is used only to bound the residual error. By contrast, the quantity C_p^{exc} below is a sum of nonnegative valuations of factors that are actually retained; it is used to prove divisibility. Cancellation in $\mathcal{E}_p$ gives no upper bound for C_p^{exc}, whose natural size can be N/p.

We therefore estimate the nonnegative valuation independently. Let

$$C_p^{\mathrm{exc}} := \sum_{\substack{a\in E,\ X_{R_y(a)}<X_0 \\ a \text{ not a bank donor}}} v_p(a) \qquad (p \leq y).$$

If $a = Rb$ occurs, then $b < 2X_0$ and $v_p(a) = v_p(b)$. For fixed b, discard the lower cutoff on R and apply the interval Selberg sieve to

$$2n/b < R \leq (2n+h)/b.$$

Its length is $h/b \gg n^{1-\delta_*}/L$. We use Selberg weights supported on $d \le y^2$; the elementary interval remainders have lcm moduli $[d,e] \le y^4 = n^{8/9}$, not merely y^2. Since $\delta_* < 1/18$, these remainders are uniformly negligible compared with the interval length after logarithmic factors. By (2.11),

$$\#\{R : 2n/b < R \le (2n+h)/b,\ P^-(R) > y\} \ll \frac{h}{b \log y} + \frac{y^4}{(\log y)^2}. \tag{6.46}$$

Now

$$\sum_{b<2X_0} \frac{v_p(b)}{b} = \sum_{k\ge1} \frac{1}{p^k} \sum_{m<2X_0/p^k} \frac{1}{m} \ll \frac{\delta_* L + 1}{p},$$

and $\sum_{b<2X_0} v_p(b) \ll X_0/p$. Multiplying (6.46) by $v_p(b)$ and summing b therefore yields, uniformly in $p \le y$, the following bound. Here $\delta_* > 0$ has already been fixed; hence, after increasing the threshold for n, one has $\delta_* L + 1 \le 2\delta_* L$. The factor two is absorbed into a fixed A_{exc}, while the displayed little-oh term below contains only the sieve remainder. Thus

$$\boxed{C_p^{\mathrm{exc}} \le A_{\mathrm{exc}} \frac{\delta_*}{\theta} \frac{N}{p} + o\Big(\frac{N}{pL}\Big)}. \tag{6.47}$$

For $p > 2X_0$, the left side is exactly zero. The sieve remainder is indeed harmless, since

$$\frac{y^4 X_0}{p(\log y)^2} = o\Big(\frac{N}{pL}\Big)$$

under $8/9 + \delta_* < 1$.

Let B_{exc} denote the product of the exceptional factors in G_{fix}. The bank and guard construction gives

$$v_p(B_0/B_{\mathrm{exc}}) = v_p(B^0_{\mathrm{bank}}) = o_W\Big(\frac{N}{pL}\Big) \qquad (p \le y), \tag{6.48}$$

by (5.18). There is no unenumerated fixed class in this estimate. Legendre's formula gives uniformly in this range

$$v_p(T) = \frac{h}{p-1} + O\Big(\frac{L}{\log p}\Big) = \frac{(c+o(1))N}{p-1}. \tag{6.49}$$

Lemma 6.8 (Combined charge)**.** *With the one-time choice* (6.1)*, the complete fixed base and the modified anchor divisor satisfy, simultaneously,*

$$D'B_0 \mid T, \qquad B_0 \mid Y := T/D'. \tag{6.50}$$

Proof. We verify the valuation inequality in three disjoint prime ranges.

If $p \in F_{\mathrm{anc}}$, the definition of γ above (6.1) gives

$$v_p(T) - v_p(D') \ge \gamma N/p.$$

Equations (6.47) and (6.48), together with $A_{\mathrm{exc}}\delta_*/\theta < \gamma/4$, show that $v_p(B_0) < \gamma N/(2p)$ for all sufficiently large n. Thus the required inequality has a fixed positive margin.

If $p \le y$ but $p \notin F_{\mathrm{anc}}$, then $v_p(D') = 0$, while (6.49) and the definition of γ give $v_p(T) \ge \gamma N/p$. The same two charge estimates again give $v_p(B_0) < \gamma N/(2p)$. This includes $p \in F_{\mathrm{hd}} \setminus F_{\mathrm{anc}}$: no anchor valuation is charged at such a prime.

Finally let $p > y$. Then $v_p(D') = 0$, because $F_{\mathrm{anc}} \subset [2, W]$. The injection (6.41) maps every factor of $G_{\mathrm{fix}} \cup G^0_{\mathrm{bank}}$ to a distinct actual factor of E, preserving the entire complete rough signature and hence every p-adic valuation. A fixed exceptional factor is mapped to itself; a bank-base factor is mapped to its designated donor. These two image classes are disjoint by construction. Therefore $v_p(B_0) \le v_p(T)$ for every $p > y$.

The three ranges cover every prime and prove $v_p(D'B_0) \le v_p(T)$ coordinatewise. This proves both assertions in (6.50). In particular, the conclusion is not obtained by multiplying two separately known divisors of T; it is a simultaneous verification for the product $D'B_0$. □

6.6 Output of the rough stage

Let $\mathcal{A} = \coprod_R \mathcal{A}_R$ be the remaining flexible lower candidates, and use the residual convention

$$r_p := v_p(Y) - v_p(B_0) - \sum_{a\in\mathcal{A}} x_a v_p(a). \tag{6.51}$$

We now give the full sign ledger rather than absorbing guard changes into words. Let $\mathcal{D}_{\rm ne}$ be the set of raw lower coordinates deleted in nonexceptional rows, and put

$$E_{\rm donor}^{\rm exc} := E_{\rm donor} \cap E_{\rm exc}.$$

For $W < p \le y$, define

$$\begin{aligned} \mathcal{R}_p^{\rm raw} &:= \sum_a \{\mathbf{1}_E(a) - x_a^{\rm raw}\} v_p(a), \\ \Delta_p^{\rm guard} &:= \sum_{a\in\mathcal{D}_{\rm ne}} x_a^{\rm raw} v_p(a) + \sum_{a\in E_{\rm donor}^{\rm exc}} v_p(a) - \sum_{a\in G_{\rm bank}^0} v_p(a). \end{aligned} \tag{6.52}$$

The first sum in $\Delta_p^{\rm guard}$ restores the raw coordinates which were removed outside the exceptional rows; the second restores the exceptional donor terms included in $\mathcal{E}_p$; and the last charges the bank base which replaced those donor tokens. Since $v_p(D') = 0$ in this medium range, direct expansion of $Y = T/D'$, $B_0 = G_{\rm fix} G_{\rm bank}^0$, and $G_{\rm fix} = E_{\rm exc} \setminus E_{\rm donor}$ gives the exact identity

$$\boxed{r_p = \mathcal{R}_p^{\rm raw} - \mathcal{E}_p - \sum_{\substack{R\ne 1\\ X_R\ge X_0}} \sum_a \nu_R(a) v_p(a) + \Delta_p^{\rm guard}.} \tag{6.53}$$

This also verifies the convention “target minus current” term by term.

For completeness, the guard census makes the last error quantitative. Prefix, row-zero, and promoted nonsmooth anchors have zero p-valuation when $W < p \le y$. For a fixed p, at most one promoted smooth anchor has nonzero p-valuation, of size at most $L/\log p$. All other terms in (6.52) come from $O(yL)$ bank endpoints or donors, and each has valuation at most $O(L/\log p)$. Hence, uniformly in the medium range,

$$|\Delta_p^{\rm guard}| \ll \frac{yL^2}{\log p} + \frac{L}{\log p} = o\left(\frac{N}{pL}\right), \tag{6.54}$$

because $pyL^3/N \le y^2L^3/N = n^{-5/9}L^4 \to 0$. Equations (6.6), (6.33), (6.36), and (6.54), inserted with their displayed signs in (6.53), prove

$$|r_p| \ll_W \frac{N}{pL} \qquad (W < p \le y). \tag{6.55}$$

Equation (6.45) proves exact equality in every nontrivial complete rough-signature row and hence $r_p = 0$ for every $p > y$, with all multiplicities. At this stage the smooth row and the coordinates $p \le W$ are provisional. The smooth-row decomposition has a frozen $F_{\rm hd}$-free protected floor and an $F_{\rm hd}$-free active excess of mass $\asymp_W N$; all fixed-head valuations will be inserted by the structured cells of the next two sections. Thus the distinction $F_{\rm anc} \subset F_{\rm hd}$, all fixed charges, and every guard are already present in the ledger to which the bridge is applied.

For reference, the exact invariants carried through the remaining stages are summarized below. “Exact rows” always means complete rough signatures, with prime-power multiplicity; “mass” includes integral fixed tokens with weight one and fractional coordinates with their actual weights.

Stage	Rough rows and $p > y$	Head $p \le W$	Medium $W < p \le y$	Total mass and ordinary log	Coordinates
After rough selection	every $R \neq 1$ exact; hence $p > y$ exact	provisional	residual bounded by (6.55)	provisional	two-sided interior on clean pools
After smooth bridge	every row exact	exact	band sums exact; pointwise residual retains the tangent bound	mass and ordinary log exact	interior
After tangent	every row exact	exact	every individual valuation exact	mass and ordinary log exact	in $[0, 1]$, with endpoint slack where used
After floating rounding	every row exact, so no $p > y$ error	integer error e_p	integer error e_p	mass exact; log changes by $\sum e_p \log p$	flexible coordinates are 0-1
After bank replacement	every row exact	exact	exact	mass and ordinary log exact	all selected factors are 0-1 and distinct

Each row of this table is proved at the indicated stage: the present section proves the first, the bridge and tangent sections prove the next two, and Lemma 10.1 and proposition 10.2 prove the last two. In particular, ordinary logarithm after rounding is not asserted separately from the valuation error; it becomes exact only when the bank changes valuation by $-e$.

7 Marked friable counts on the structured cells

The smooth bridge uses only physical intervals and exact finite valuation and coprimality patterns at the head primes. This restriction is essential: ordinary smooth-number estimates do not count an arbitrary nonzero residue class modulo a fixed integer. In this section we prove all one- and two-marked estimates needed for these structured cells, including their stability under the compact tilts used later.

Throughout,

$$y = n^{2/9}, \qquad U := \frac{\log n}{\log y} = \frac{9}{2}, \qquad L = \log n.$$

The numerical value $U = 9/2$ is used below. In particular, we make no terminal prime-power assertion at $U = 5/2$.

7.1 Exact head cells and a four-mark estimate

Fix a finite set $H \subset F_{\mathrm{hd}}$, a vector of fixed nonnegative integers $\mathbf{e} = (e_\ell)_{\ell \in H}$, and fixed $0 < A < B < \infty$. Put

$$h_{\mathbf{e}} = \prod_{\ell \in H} \ell^{e_\ell}, \qquad M_H = \prod_{\ell \in H} \ell, \qquad \delta_{\mathbf{e}} = \frac{1}{h_{\mathbf{e}}} \frac{\varphi(M_H)}{M_H},$$

and define the structured cell

$$\mathcal{C}_{\mathbf{e}}(A, B) = \{m : An < m \le Bn,\ P^+(m) \le y,\ v_\ell(m) = e_\ell\ (\ell \in H)\}.$$

The finite collection of $A, B, H, \mathbf{e}$ used in the construction is fixed before $n \to \infty$. For a bounded-variation function we abbreviate $\|\phi\|_{BV} := \|\phi\|_\infty + \operatorname{Var} \phi$.

For a y-smooth integer d, write

$$v(d) = \frac{\log d}{\log y}, \qquad h_U(v) = \frac{\rho(U-v)}{\rho(U)}.$$

If ν_i is normalized uniform measure on a structured cell C_i, a *finite mixture of cells* means the probability measure

$$\nu = \sum_{i=1}^{m} \lambda_i \nu_i, \qquad \lambda_i \geq 0, \qquad \sum_{i=1}^{m} \lambda_i = 1, \tag{7.1}$$

where m is fixed. When compact positivity is required later, the weight vector is restricted to a fixed compact subset of the relative interior of the relevant simplex. A piecewise-constant baseline weight on the disjoint cells is exactly of the form (7.1) after normalization. We apply count formulas to each C_i; only probabilities, expectations, and covariances are then averaged under ν. We never interpret a nonintegral convex mixture as a count.

Proposition 7.1 (Uniform marked structured cells)**.** *Suppose that $d \leq y^4$ is y-smooth and $(d, M_H) = 1$. Uniformly in d,*

$$\begin{aligned} A_{\mathbf{e},A,B}(d) &:= \#\{m \in \mathcal{C}_{\mathbf{e}}(A,B) : d \mid m\} \\ &= \delta_{\mathbf{e}} \frac{(B-A)n}{d} \left\{ \rho(U - v(d)) + O_{H,A,B}(L^{-1}) \right\}. \end{aligned} \tag{7.2}$$

Consequently, for the normalized uniform measure on the cell,

$$\mathbb{P}_{\mathbf{e},A,B}(d \mid m) = \frac{h_U(v(d))}{d} + O_{H,A,B}\left(\frac{1}{dL}\right). \tag{7.3}$$

If $\phi : [A,B] \to \mathbb{R}$ has bounded variation, then

$$\begin{aligned} &\sum_{\substack{m \in \mathcal{C}_{\mathbf{e}}(A,B) \\ d \mid m}} \phi(m/n) \\ &= \delta_{\mathbf{e}} \frac{n}{d} \rho(U - v(d)) \int_A^B \phi(t)\, dt + O_{H,A,B}\left(\frac{n}{dL}\{\|\phi\|_\infty + \operatorname{Var}\phi\}\right). \end{aligned} \tag{7.4}$$

The normalized assertions hold for a fixed finite disjoint union, or a probability mixture in the sense of (7.1). Any overlapping family is first disjointified and then interpreted through its induced convex mixture. In particular, any two allowed normalized cells C, C' satisfy

$$|\mathbb{P}_C(d \mid m) - \mathbb{P}_{C'}(d \mid m)| \ll_W \frac{1}{dL}. \tag{7.5}$$

Proof. The exact valuation pattern, unlike a general residue class, has the finite divisibility identity

$$\mathbf{1}_{\{v_\ell(m) = e_\ell \ (\ell \in H)\}} = \mathbf{1}_{h_{\mathbf{e}} \mid m} \sum_{a \mid M_H} \mu(a) \mathbf{1}_{a \mid m/h_{\mathbf{e}}}.$$

Since $(d, M_H) = 1$, it follows exactly that

$$A_{\mathbf{e},A,B}(d) = \sum_{a \mid M_H} \mu(a) \left\{ \Psi\left(\frac{Bn}{h_{\mathbf{e}} ad}, y\right) - \Psi\left(\frac{An}{h_{\mathbf{e}} ad}, y\right) \right\}. \tag{7.6}$$

The smallest unscaled cofactor is

$$\frac{n}{y^4} = n^{1/9} = y^{1/2} \longrightarrow \infty. \tag{7.7}$$

Thus every smooth parameter in (7.6) lies in the fixed compact interval $[1/2+o(1), 9/2+o(1)]$. When it is at least one, apply Theorem 2.1 and (2.3) [HT93, Sai89]; when it is at most one, use the exact identity $\Psi(x,y) = \lfloor x \rfloor$. The fixed factors $A, B, h_{\mathbf{e}}, a$ shift the parameter by $O_H(1/L)$, and ρ is Lipschitz on this compact range. Hence, uniformly for $C = A, B$,

$$\Psi\left(\frac{Cn}{h_{\mathbf{e}}ad}, y\right) = \frac{Cn}{h_{\mathbf{e}}ad}\left\{\rho(U - v(d)) + O_H(L^{-1})\right\} + O(1). \tag{7.8}$$

By (7.7), the endpoint $O(1)$ is $O(n/(dL))$. Substitute (7.8) in (7.6) and use

$$\frac{1}{h_{\mathbf{e}}} \sum_{a|M_H} \frac{\mu(a)}{a} = \delta_{\mathbf{e}}.$$

This proves (7.2). With $d = 1$, it also gives

$$\#\mathcal{C}_{\mathbf{e}}(A,B) = \delta_{\mathbf{e}}(B-A)n\{\rho(U) + O_W(L^{-1})\} \asymp_W n, \tag{7.9}$$

so division proves (7.3).

Apply (7.2) uniformly to every cumulative subinterval $(An, tn]$, $A \le t \le B$. Its error is $O_W(n/(dL))$, independent of t. Stieltjes partial summation against ϕ proves (7.4). Every cell has the same normalized main divisibility profile, so finite disjoint unions and compact convex mixtures preserve it; this also proves (7.5). $\square$

The cutoff y^4 covers exactly the divisors used later:

$$p, \quad pq, \quad p^k, \quad p^k q, \quad p^k q r^j,$$

where $W < p, q, r \le y$, $p^k \le y^2$, and $r^j \le y$. In particular, the last product is at most y^4. This is the genuine reason for fixing $U = 9/2$: all pointwise marked products retain the cofactor margin (7.7).

7.2 One and two marks, prime powers, and physical tests

Let C be an allowed cell or finite mixture and put $t_p = \log p / \log y$. For distinct $W < p, q \le y$, Proposition 7.1 gives

$$\mathbb{E}_C \mathbf{1}_{p|m} = \frac{h_U(t_p)}{p} + O_W(1/(pL)),$$
$$\mathbb{E}_C \mathbf{1}_{pq|m} = \frac{h_U(t_p + t_q)}{pq} + O_W(1/(pqL)).$$

It follows that

$$\mathrm{Var}_C(\mathbf{1}_{p|m}) = \frac{h_U(t_p)}{p} + O_W(1/(pL) + 1/p^2),$$
$$\mathrm{Cov}_C(\mathbf{1}_{p|m}, \mathbf{1}_{q|m}) = \frac{h_U(t_p + t_q) - h_U(t_p)h_U(t_q)}{pq} + O_W(1/(pqL)). \tag{7.10}$$

Likewise, if $p^k \le y^2$ and $p \ne q$, then

$$\mathbb{E}_C \mathbf{1}_{p^k q|m} = \frac{h_U(kt_p + t_q)}{p^k q} + O_W(1/(p^k qL)),$$
$$\mathrm{Cov}_C(\mathbf{1}_{p^k|m}, \mathbf{1}_{q|m}) = \frac{h_U(kt_p + t_q) - h_U(kt_p)h_U(t_q)}{p^k q} + O_W(1/(p^k qL)).$$

There is also a comparison for the full valuation. Expand $v_p = \sum_{k\ge 1} \mathbf{1}_{p^k|m}$. Applying (7.5) for $p^k \le y^4$ gives total error $O_W(1/(pL))$. The remaining powers satisfy, by the elementary multiple count and (7.9),

$$\sum_{p^k > y^4} \mathbb{P}_C(p^k \mid m) \ll y^{-4} + L/n = o(1/(pL)).$$

Consequently

$$|\mathbb{E}_C v_p(m) - \mathbb{E}_{C'} v_p(m)| \ll_W \frac{1}{pL} \qquad (W < p \le y). \tag{7.11}$$

If ϕ has bounded variation on the common physical compact set, (7.4) and (7.3) give

$$|\operatorname{Cov}_C(\mathbf{1}_{d|m}, \phi(m/n))| \ll_W \frac{\|\phi\|_{BV}}{dL}. \tag{7.12}$$

The same estimate holds for finite mixtures: within-cell covariances have this bound, while between-cell covariances are controlled by (7.5). Summing over $d = p^k$ yields

$$|\operatorname{Cov}_C(v_p(m), \phi(m/n))| \ll_W \frac{\|\phi\|_{BV}}{pL}.$$

For example, with $\phi(t) = \log t$ and arbitrary coefficients c_p,

$$\left|\operatorname{Cov}_C\left(\sum_{W<p\le y} c_p v_p(m), \log(m/n)\right)\right| \ll_W \frac{1}{L} \sum_{W<p\le y} \frac{|c_p|}{p}.$$

7.3 Stability under compact bridge tilts

The bridge changes the measure by bounded band fugacities divided by L. We prove that the error scale above survives this homotopy. Let

$$S(m) = \sum_{W<p\le y} \eta_p v_p(m), \qquad |\eta_p| \le B, \tag{7.13}$$

where B is fixed, and denote expectation after reweighting by $e^{S/L}$ by $\mathbb{E}_{C,S}$.

Lemma 7.2 (Squarefree marked transfer under tilts)**.** *Uniformly for $d = p$ or $d = pq$, with distinct primes in $(W, y]$,*

$$\mathbb{P}_{C,S}(d \mid m) = \mathbb{P}_C(d \mid m) + O_{B,W}(1/(dL)). \tag{7.14}$$

If ϕ is a fixed bounded-variation physical test, then

$$|\operatorname{Cov}_{C,S}(\mathbf{1}_{d|m}, \phi(m/n))| \ll_{B,W} \frac{\|\phi\|_{BV}}{dL}. \tag{7.15}$$

Proof. Put $\Omega(m) = \sum_{W<p\le y} v_p(m)$. Since $\Omega(m)\log W \le \log(Bn)$, the factor $e^{B\Omega/L}$ is uniformly bounded. Moreover, for $j = 1, 2$, elementary counting gives

$$\mathbb{E}_C\{\mathbf{1}_{d|m}\Omega(m)^j\} \ll_{W,j} \frac{(1+\log L)^j}{d}. \tag{7.16}$$

Indeed write $m = dk$, discard smoothness and the head restrictions, and enlarge to $k \le Bn/d$. Since d has at most two prime factors, $\Omega(dk) \le 2 + \Omega(k)$. Expanding the first two powers into prime-power divisibility indicators and using $\lfloor X/e \rfloor \le X/e$ gives

$$\frac{1}{X}\sum_{k\le X} \Omega(k)^j \ll_j \left(1 + \sum_{W<r\le y}\sum_{a\ge 1} \frac{1}{r^a}\right)^j \ll_j (1+\log L)^j.$$

Division by the cell size (7.9) proves (7.16).

Taylor's formula, with the uniform exponential bound just noted, gives

$$\mathbb{E}_C e^{S/L} = 1 + L^{-1}\mathbb{E}_C S + O_B((\log L)^2/L^2),$$
$$\mathbb{E}_C(\mathbf{1}_{d|m} e^{S/L}) = \mathbb{P}_C(d \mid m) + L^{-1}\mathbb{E}_C(\mathbf{1}_{d|m} S) + O_B((\log L)^2/(dL^2)). \tag{7.17}$$

It remains to use cancellation in the centered first-order term. Expand S into prime-power indicators. If the score prime r is not a factor of d, Proposition 7.1 applies to dr^a for $r^a \le y$; the product is then at most y^3. Its main covariance is bounded by

$$\frac{1}{dr^a}|h_U(v(d) + at_r) - h_U(v(d))h_U(at_r)| \ll_U \frac{at_r}{dr^a},$$

because the expression vanishes at $at_r = 0$ and h_U is Lipschitz on the relevant compact range. The at most two score primes dividing d contribute $O_B(1/d)$ directly. The marked errors and the tail $r^a > y$ are summable by elementary common-multiple bounds. Since

$$\sum_{W<r\le y}\sum_{r^a\le y} \frac{at_r}{r^a} = O_W(1),$$

we obtain

$$|\operatorname{Cov}_C(\mathbf{1}_{d|m}, S)| \ll_{B,W} \frac{1}{d}.$$

Divide the two lines of (7.17) and use $(\log L)^2/L^2 = o(1/L)$. This proves (7.14).

Insert $\phi(m/n)$ into the same calculation. In every leading marked term, (7.4) factors the same Lebesgue average of ϕ. More explicitly, the first derivative at zero of the tilted covariance is the third centered combination

$$\begin{aligned}\kappa_C(\mathbf{1}_d, \phi, S) := \mathbb{E}_C(\mathbf{1}_d\phi S) - \mathbb{E}_C(\mathbf{1}_d\phi)\mathbb{E}_C S - \mathbb{E}_C(\mathbf{1}_d S)\mathbb{E}_C\phi \\ - \mathbb{E}_C(\phi S)\mathbb{P}_C(d \mid m) + 2\mathbb{P}_C(d \mid m)\mathbb{E}_C\phi\,\mathbb{E}_C S.\end{aligned}$$

For each prime-power term of S, the common Stieltjes main term in these five summands cancels exactly. Summing the remaining Stieltjes errors gives

$$\frac{1}{L}|\kappa_C(\mathbf{1}_d, \phi, S)| \ll_{B,W} \frac{\|\phi\|_{BV}\log L}{dL^2}.$$

The three-factor Taylor remainder is $O_{B,W}(\|\phi\|_{BV}(\log L)^2/(dL^2))$, hence is $o(\|\phi\|_{BV}/(dL))$. Together with the un-tilted covariance (7.12), this proves (7.15). □

Individual very high prime powers need not satisfy a relative $O(1/L)$ stability statement when the tilt itself contains that prime. The bridge does not need such a statement. It needs the following aggregate transfer, which we now prove.

For a set E of at most two primes, let S_E be the score (7.13) with the primes of E omitted. The extension from squarefree marks to arbitrary powers supported on E is recorded separately because it is used in every local restoration.

Lemma 7.3 (Omitted-local-score transfer)**.** *If $d \le y^3$ is supported on E, then, uniformly in d,*

$$\frac{\mathbb{E}_C(\mathbf{1}_{d|m} e^{S_E/L})}{\mathbb{E}_C e^{S_E/L}} = \frac{h_U(v(d))}{d} + O_{B,W}(1/(dL)). \tag{7.18}$$

For two divisibility indicators the same assertion holds with d replaced by their least common multiple, provided that lcm is at most y^3. If a bounded-variation physical test $\phi(m/n)$ is inserted, its leading Lebesgue mean factors and the remaining error is $O_{B,W}(\|\phi\|_{BV}/(dL))$.

Proof. Put

$$\Omega_E(m) = \sum_{\substack{W<r\le y\\ r\notin E}} v_r(m).$$

Writing $m = dk$, discarding smoothness and the fixed cell restrictions, and expanding one or two powers of Ω_E into divisibility indicators gives, exactly as an elementary multiple-counting estimate,

$$\mathbb{E}_C\{\mathbf{1}_{d|m}\Omega_E(m)^j\} \ll_{W,j} \frac{(1+\log L)^j}{d} \qquad (j=1,2). \tag{7.19}$$

The endpoint contribution is absorbed because $d \le y^3$ and $n/y^3 = n^{1/3}$. Since $|S_E| \le B\Omega_E$ and $e^{B\Omega_E/L} \ll_{B,W} 1$, Taylor's formula reduces the desired ratio to bounding the centered first-order term.

Expand S_E as $\sum_{r\notin E}\eta_r \sum_{a\ge 1}\mathbf{1}_{r^a|m}$. When $r^a \le y$, the product $dr^a \le y^4$, and Proposition 7.1 gives the main covariance

$$\frac{h_U(v(d)+at_r) - h_U(v(d))h_U(at_r)}{dr^a}.$$

The numerator vanishes at $at_r = 0$ and is $O_U(at_r)$ by Lipschitz continuity. Hence these main terms sum to $O_{B,W}(1/d)$, because

$$\sum_{W<r\le y}\sum_{r^a\le y}\frac{at_r}{r^a} = O_W(1).$$

Their marked errors sum to $O_{B,W}((1+\log L)/(dL))$. For $r^a > y$, elementary common-multiple counting gives

$$\mathbb{P}_C(dr^a \mid m) \ll_W \frac{1}{dr^a} + \frac{1}{n}.$$

The reciprocal tails are $O_W(1/d)$; the literal endpoints total $O_W(yL/n) = O_W(1/d)$, since $d \le y^3$ and $y^4L = o(n)$. We have therefore proved

$$|\operatorname{Cov}_C(\mathbf{1}_{d|m}, S_E)| \ll_{B,W} \frac{1}{d}.$$

Using (7.19), the second-order Taylor remainder is $O_{B,W}((1+\log L)^2/(dL^2)) = o(1/(dL))$. Division by the analogous denominator expansion, followed by (7.3), proves (7.18). A product of two divisibility indicators is the indicator of their lcm. Finally, inserting ϕ and using (7.4) makes the same Lebesgue average appear in every leading centered term; it cancels, while the Stieltjes errors are bounded by $O_{B,W}(\|\phi\|_{BV}/(dL))$. This proves all assertions. □

Restore an omitted prime p. If $\lambda_p = e^{\eta_p/L}$, the exact identity

$$\lambda_p^{v_p(m)} = 1 + \sum_{a\ge 1}(\lambda_p^a - \lambda_p^{a-1})\mathbf{1}_{p^a|m} \tag{7.20}$$

has

$$|\lambda_p^a - \lambda_p^{a-1}| \le \frac{C_B}{L}e^{Ba/L}. \tag{7.21}$$

Let $C_* > 0$ be a fixed upper bound for every physical endpoint in the finite cell family. Fix once and for all a bounded Lipschitz extension $F : [0,\infty) \to \mathbb{R}$ of $h_U|_{[0,3]}$, with $F(0) = 1$. For local fugacities $\lambda_p = e^{\eta_p/L}$, define the finite coefficients

$$A_p = \left\lfloor \frac{\log(C_* n)}{\log p} \right\rfloor, \qquad c_{p,0} = 1, \qquad c_{p,a} = \begin{cases} \lambda_p^a - \lambda_p^{a-1}, & 1 \le a \le A_p, \\ 0, & a > A_p, \end{cases} \tag{7.22}$$

and similarly for q. The explicit two-prime model is

$$Q_{p,q}(r,s) := \sum_{a,b\geq 0} c_{p,a}c_{q,b}\frac{F(\max(r,a)t_p+\max(s,b)t_q)}{p^{\max(r,a)}q^{\max(s,b)}}. \tag{7.23}$$

The following lemma supplies the two-local-factor calculation, including the terms beyond the four-mark range.

Lemma 7.4 (Two-local-factor expansion)**.** *Let $p \neq q$ lie in $(W,y]$, put $I_{p^r} = \mathbf{1}_{p^r|m}$, and allow $r,s\geq 0$, with $I_{p^0} = I_{q^0} = 1$. If $p^rq^s \leq y^3$, then the exact restoration of the two omitted local factors satisfies the explicit formula*

$$\mathbb{E}_{C,S}(I_{p^r}I_{q^s}) = \frac{Q_{p,q}(r,s)}{Q_{p,q}(0,0)} + O_{B,W}\left(\frac{1}{p^rq^sL}+\frac{1}{n}\right).$$

Moreover $Q_{p,q}(0,0) = 1 + O_{B,W}(1/(Lp)+1/(Lq)+1/n)$, so the denominator is bounded away from zero for all sufficiently large n. In particular, whenever $p^k \leq y^2$,

$$|\operatorname{Cov}_{C,S}(I_{p^k}, I_q)| \leq C_B\frac{t_q}{p^kq} + O_{B,W}\left(\frac{1}{p^kqL}\right). \tag{7.24}$$

The transposed assertion is

$$|\operatorname{Cov}_{C,S}(I_p, I_{q^l})| \leq C_B\frac{t_p}{pq^l} + O_{B,W}\left(\frac{1}{pq^lL}\right) \qquad (q^l \leq y^2). \tag{7.25}$$

Proof. Let ν be the probability law obtained by tilting the allowed base law ν_C by $e^{S_{\{p,q\}}/L}$. If C is an actual cell, ν_C is its normalized counting measure; if C denotes a finite mixture, ν_C is the probability mixture defined in (7.1). The density of ν relative to ν_C is bounded above and below by constants depending only on B,W. Elementary multiple counting and (7.9) give the following estimate on every actual cell. For a mixture, apply that estimate cellwise to the tilted numerator and use the uniform two-sided bounds for the (cellwise and mixture) normalizing constants. Thus, for every integer D,

$$\nu(D \mid m) \ll_{B,W} \frac{1}{D}+\frac{1}{n}. \tag{7.26}$$

For $D \leq y^3$ supported on p,q, the stronger formula (7.18) applies.

One local factor. Use the coefficients from (7.22). Since $v_p(m) \leq A_p$ throughout the cell family, this truncation leaves (7.20) exact. From (7.21), geometric summation gives, uniformly for $r \geq 0$,

$$\sum_{a\geq 1}|c_{p,a}| \ll_{B,W} 1, \qquad \sum_{a\geq 1}\frac{|c_{p,a}|}{p^a} \ll_B \frac{1}{Lp}, \qquad \sum_{a\geq 0}\frac{|c_{p,a}|}{p^{\max(r,a)}} \ll_{B,W} \frac{1}{p^r}. \tag{7.27}$$

The first estimate also follows directly by telescoping separately when $\lambda_p \geq 1$ and $\lambda_p < 1$; indeed $A_p = O_W(L)$.

Two local factors. Multiplying the two identities (7.20) gives the exact numerator

$$\sum_{a,b\geq 0} c_{p,a}c_{q,b}\,\nu\Big(p^{\max(r,a)}q^{\max(s,b)} \mid m\Big).$$

The denominator is the same expression with $r = s = 0$. Split the (a,b)-sum according as

$$p^{\max(r,a)}q^{\max(s,b)} \leq y^3 \quad \text{or} \quad p^{\max(r,a)}q^{\max(s,b)} > y^3. \tag{7.28}$$

On the first part insert (7.18). By (7.27), the sum of all its marked errors is

$$\ll_{B,W} \frac{1}{L} \sum_{a,b\geq 0} \frac{|c_{p,a}c_{q,b}|}{p^{\max(r,a)}q^{\max(s,b)}} \ll_{B,W} \frac{1}{p^r q^s L}.$$

On the complementary part, at least one of $a > r$ or $b > s$ holds. Using (7.21) in that coordinate and (7.27) in the other gives explicitly

$$\begin{aligned}&\sum_{\substack{a,b\geq 0\\ p^{\max(r,a)}q^{\max(s,b)}>y^3}} \frac{|c_{p,a}c_{q,b}|}{p^{\max(r,a)}q^{\max(s,b)}}\\ &\quad\leq \left(\sum_{a>r}\frac{|c_{p,a}|}{p^a}\right)\left(\sum_{b\geq 0}\frac{|c_{q,b}|}{q^{\max(s,b)}}\right) + \left(\sum_{a\geq 0}\frac{|c_{p,a}|}{p^{\max(r,a)}}\right)\left(\sum_{b>s}\frac{|c_{q,b}|}{q^b}\right) \ll_{B,W} \frac{1}{p^r q^s L}.\end{aligned} \tag{7.29}$$

The endpoint part of (7.26) is separate and is

$$\frac{1}{n}\sum_{a,b\geq 0}|c_{p,a}c_{q,b}| \ll_{B,W} \frac{1}{n}.$$

This proves the asserted tail split with every floor contribution retained.

For completeness, we now identify the covariance in the explicit model (7.23). Replacing the terms outside (7.28) by this extension costs no more than (7.29). Hence

$$\mathbb{E}_{C,S}(I_{p^r}I_{q^s}) = \frac{Q_{p,q}(r,s)}{Q_{p,q}(0,0)} + O_{B,W}\left(\frac{1}{p^r q^s L} + \frac{1}{n}\right), \tag{7.30}$$

The $(a,b) = (0,0)$ term of $Q_{p,q}(0,0)$ is one, while (7.27) bounds all remaining terms by $O_{B,W}(1/(Lp) + 1/(Lq))$. This proves the denominator assertion in the lemma (with harmless room for the literal endpoint), and in particular $Q_{p,q}(0,0) \asymp_{B,W} 1$.

Rank-one cancellation. Write

$$K_F(x,z) = F(x+z) - F(x)F(z).$$

Lipschitz continuity and $F(0) = 1$ imply

$$|K_F(x,z)| \ll_U z, \qquad K_F(0,z) = K_F(x,0) = 0. \tag{7.31}$$

If

$$\mathcal{P}_p(r) = \sum_{a\geq 0} c_{p,a}\frac{F(\max(r,a)t_p)}{p^{\max(r,a)}},$$

and $\mathcal{P}_q(s)$ is defined similarly, then $Q_{p,q}(r,s) = \mathcal{P}_p(r)\mathcal{P}_q(s) + \mathcal{R}(r,s)$. Equations (7.27) and (7.31) give

$$\begin{aligned}
|\mathcal{P}_p(r)| &\ll_{B,W} p^{-r}, & |\mathcal{P}_q(1)| &\ll_{B,W} q^{-1},\\
|\mathcal{R}(r,1)| &\ll_{B,W} \frac{t_q}{p^r q}, & |\mathcal{R}(r,0)| &\ll_{B,W} \frac{t_q}{p^r Lq},\\
|\mathcal{R}(0,1)| &\ll_{B,W} \frac{t_q}{Lpq}, & |\mathcal{R}(0,0)| &\ll_{B,W} \frac{t_q}{L^2 pq}.
\end{aligned} \tag{7.32}$$

For example, the q-sum in the first remainder is bounded by

$$\sum_{b\geq 0}|c_{q,b}|\frac{\max(1,b)t_q}{q^{\max(1,b)}} \leq \frac{t_q}{q} + \frac{C_B t_q}{L}\sum_{b\geq 1}\frac{b}{q^b} \ll_B \frac{t_q}{q},$$

whereas for $\mathcal{R}(r,0)$ its $b=0$ term vanishes and the extra factor L^{-1} remains. The two analogous p-sums give the last two bounds.

Expand $Q_{p,q}(r,1)Q_{p,q}(0,0)-Q_{p,q}(r,0)Q_{p,q}(0,1)$. The rank-one products $\mathcal{P}_p\mathcal{P}_q$ cancel exactly, and (7.32) bounds what remains by

$$|Q_{p,q}(r,1)Q_{p,q}(0,0)-Q_{p,q}(r,0)Q_{p,q}(0,1)| \ll_{B,W} \frac{t_q}{p^r q}.$$

Combining this with the four instances of (7.30) gives

$$|\operatorname{Cov}_{C,S}(I_{p^r},I_q)| \le C_B \frac{t_q}{p^r q} + O_{B,W}\left(\frac{1}{p^r qL}+\frac{1}{n}\right).$$

When $p^r \le y^2$, the endpoint term is absorbed because $p^r q \le y^3$ and $n/y^3 = n^{1/3}$. This proves (7.24). Interchanging p,r and q,s proves (7.25). $\square$

We now perform all four aggregate prime-power sums. Let

$$I_p = \mathbf{1}_{p|m}, \qquad V_p = v_p(m), \qquad J_p := V_p - I_p = \sum_{k\ge 2} I_{p^k},$$

and use the weighted row norm

$$\|K\|_{\mathrm{row}} := \sup_{W<p\le y} p \sum_{W<q\le y} |K(p,q)|. \tag{7.33}$$

The order in which the local powers are summed matters for the later moving-low-cell argument. We therefore record a box-uniform refinement of the preceding two-local calculation. Its leading constant is independent of the radius of the compact tilt box.

Lemma 7.5 (Product-weighted aggregate power transfer)**.** *There is a constant $C_{\mathrm{pow}}>0$, depending only on the fixed Dickman parameter and the fixed physical compact set, with the following property. For every fixed B,W there is a quantity $\epsilon_{B,W}(n)\to 0$ such that, for distinct primes $W<p,q\le y$,*

$$\sum_{k\ge 2} |\operatorname{Cov}_{C,S}(I_{p^k},I_q)| \le C_{\mathrm{pow}}\frac{t_p t_q}{p^2 q} + \frac{\epsilon_{B,W}(n)}{p^2 q} + O_{B,W}(L/n), \tag{7.34}$$

$$\sum_{l\ge 2} |\operatorname{Cov}_{C,S}(I_p,I_{q^l})| \le C_{\mathrm{pow}}\frac{t_p t_q}{p q^2} + \frac{\epsilon_{B,W}(n)}{p q^2} + O_{B,W}(L/n),$$

$$\sum_{k,l\ge 2} |\operatorname{Cov}_{C,S}(I_{p^k},I_{q^l})| \le C_{\mathrm{pow}}\frac{t_p t_q}{p^2 q^2} + \frac{\epsilon_{B,W}(n)}{p^2 q^2} + O_{B,W}(L^2/n). \tag{7.35}$$

Moreover

$$\mathbb{E}_{C,S} J_p^2 \le \frac{C_{\mathrm{pow}}}{p^2} + \frac{\epsilon_{B,W}(n)}{p^2} + O_{B,W}(L^2/n). \tag{7.36}$$

Consequently

$$\|\operatorname{Cov}_{C,S}(V_p,V_q) - \operatorname{Cov}_{C,S}(I_p,I_q)\|_{\mathrm{row}} \le \frac{C_{\mathrm{pow}}}{W} + \epsilon_{B,W}(n), \tag{7.37}$$

after enlarging C_{pow} by an absolute factor. In particular, the coefficient of $1/W$ in this estimate does not depend on B. Here $\epsilon_{B,W}(n)$ denotes one fixed maximum of the finitely many pair, tail, endpoint, and row remainders appearing in the proof; the explicit admissible choice is given at the end of the proof. Thus the same function is used simultaneously in all four displays.

Proof. Four-mark chamber transfer. We first transfer the structured-cell formula through the complete four-mark range while omitting the forced local primes from the tilt. For a set E of one or two primes in $(W, y]$, put

$$\Omega_E(m) = \sum_{\substack{W<r\le y\\ r\notin E}} v_r(m), \qquad S_E(m) = \sum_{\substack{W<r\le y\\ r\notin E}} \eta_r v_r(m),$$

and let ν_E be the probability law obtained by tilting the allowed base law ν_C by $e^{S_E/L}$. If D is supported on E, then $\Omega_E(Da) = \Omega_E(a)$. For each actual cell C_i, one has $\#C_i \asymp_W n$, and elementary multiple counting consequently gives the following bound on C_i, uniformly for $D \le y^4$. Since the number of cells is fixed, averaging the cellwise expectations with the weights in (7.1) gives exactly the same bound for an allowed finite mixture:

$$\mathbb{E}_C\Big(\mathbf{1}_{D|m}\Omega_E(m)^j\Big) \ll_{W,j} \frac{(1+\log L)^j}{D} \qquad (j=1,2). \tag{7.38}$$

Here are the details needed for the second moment. On writing $m = Da$ and discarding smoothness and the cell restrictions, expand

$$\Omega_E(a) = \sum_{\substack{W<r\le y\\ r\notin E}} \sum_{b\ge 1} \mathbf{1}_{r^b|a}.$$

The distinct-prime terms in its square have mean at most A_y^2, where

$$A_y := \sum_{W<r\le y}\sum_{b\ge 1}\frac{1}{r^b} = \sum_{W<r\le y}\frac{1}{r-1} \ll_W 1+\log L.$$

The same-prime terms have mean at most

$$\sum_{W<r\le y}\sum_{a,b\ge 1}\frac{1}{r^{\max(a,b)}} = \sum_{W<r\le y}\sum_{c\ge 1}\frac{2c-1}{r^c} \ll_W 1+\log L.$$

This proves (7.38) literally.

We have $|S_E| \le B\Omega_E$, while $\Omega_E(m)\log W \le \log(C_* n)$ for a fixed physical endpoint C_*. Thus $e^{B\Omega_E/L} \ll_{B,W} 1$ pointwise. Taylor's formula and (7.38) show that

$$\mathbb{E}_C(\mathbf{1}_{D|m}e^{S_E/L}) = \mathbb{P}_C(D \mid m) + O_{B,W}\left(\frac{1+\log L}{DL} + \frac{(1+\log L)^2}{DL^2}\right),$$
$$\mathbb{E}_C e^{S_E/L} = 1 + O_{B,W}\left(\frac{1+\log L}{L} + \frac{(1+\log L)^2}{L^2}\right).$$

The marked formula (7.3), valid uniformly through $D \le y^4$, may now be inserted before division. We obtain

$$\nu_E(D \mid m) = \frac{h_U(v(D))}{D} + O_{B,W}\left(\frac{1+\log L}{DL}\right) \qquad (D \le y^4, \operatorname{supp} D \subset E). \tag{7.39}$$

There is no separate endpoint here: $n/D \ge n^{1/9}$, so the endpoint in the structured-cell count is absorbed by $1/(DL)$. For arbitrary D, the bounded density ratio and elementary multiple counting give the fallback

$$\nu_E(D \mid m) \ll_{B,W} D^{-1} + n^{-1}. \tag{7.40}$$

We next record the product estimate for the Dickman profile using only the regularity available across its kink. Put $H = h_U|_{[0,4]}$. The function H is bounded and Lipschitz, $H(0) = 1$, and it is $C^{1,1}$ on a fixed neighborhood of zero. Hence

$$|H(x+z) - H(x)H(z)| \le C_K xz \qquad (x, z \ge 0,\ x+z \le 4), \tag{7.41}$$

where C_K depends only on U. Indeed, in a sufficiently small square at the origin the kernel vanishes on both axes and its mixed derivative is bounded. If x is in that small interval and z is outside it, Lipschitz continuity and $H(0) = 1$ give

$$|H(x+z) - H(x)H(z)| \leq |H(x+z) - H(z)| + |H(z)|\,|H(x) - 1| \ll x \ll xz;$$

the transposed case is identical. When both variables are outside the small interval, boundedness gives $O(1) \ll xz$. This proves (7.41); in particular, no $C^{1,1}$ assertion on all of $[0,4]$ is being used.

It remains to restore the primes of E exactly. Use the truncated coefficients $c_{p,a}$ from the proof of Lemma 7.4, so that (7.20) is exact on every allowed cell. Fix B, W and take n large enough that $\lambda_p \leq 3/2$. For $r \geq 0$, define

$$D_p(r) := \sum_{a \geq 1} \frac{|c_{p,a}|}{p^{\max(r,a)}}, \qquad C_p(r) := p^{-r} + D_p(r).$$

Splitting at $a = r$, using $|c_{p,a}| \leq (2B/L)\lambda_p^{a-1}$ and $\lambda_p^{A_p} \ll_{B,W} 1$, and summing the geometric tail gives

$$\sum_{a \geq 1} |c_{p,a}| \ll_{B,W} 1, \qquad D_p(0) \ll_B \frac{1}{Lp},$$
$$D_p(r) \ll_{B,W} \frac{r+1}{Lp^r} \quad (r \geq 1), \qquad C_p(r) \ll_{B,W} p^{-r} \quad (r \geq 1),$$

and therefore

$$\sum_{r \geq 2} D_p(r) \ll_{B,W} \frac{1}{Lp^2}, \qquad \sum_{r \geq 2} C_p(r) \ll_{B,W} \frac{1}{p^2},$$
$$\sum_{r \geq 2} (2r-3) D_p(r) \ll_{B,W} \frac{1}{Lp^2}, \qquad \sum_{r \geq 2} \frac{2r-3}{p^r} \ll \frac{1}{p^2}. \tag{7.42}$$

For example, in $D_p(r)$ the two ranges satisfy

$$p^{-r} \sum_{1 \leq a \leq r} |c_{p,a}| \ll_{B,W} \frac{r}{Lp^r}, \qquad \sum_{a > r} \frac{|c_{p,a}|}{p^a} \ll_{B,W} \frac{1}{Lp^{r+1}}.$$

These estimates also prove the aggregate displays.

Take first $E = \{p, q\}$. If $N_{r,s}$ is the numerator for $p^r q^s \mid m$ after restoring the two local fugacities, exact expansion gives

$$N_{r,s} = \sum_{a,b \geq 0} c_{p,a} c_{q,b}\, \nu_E\Big(p^{\max(r,a)} q^{\max(s,b)} \mid m\Big), \qquad \mathbb{P}_{C,S}(p^r q^s \mid m) = \frac{N_{r,s}}{N_{0,0}}.$$

Apply (7.40) to every term other than $(a,b) = (0,0)$. The coefficient sums just proved yield

$$|N_{r,s} - \nu_E(p^r q^s \mid m)| \ll_{B,W} D_p(r) C_q(s) + p^{-r} D_q(s) + n^{-1}. \tag{7.43}$$

The endpoint is really $O_{B,W}(n^{-1})$, because the unweighted absolute coefficient sums are bounded. Taking $r = s = 0$ shows

$$N_{0,0} = 1 + O_{B,W}\left(\frac{1}{Lp} + \frac{1}{Lq} + \frac{1}{n}\right). \tag{7.44}$$

Thus the local denominator is $1 + o_{B,W}(1)$, not merely bounded by a box-dependent constant.

For clarity, we make the summation implicit in the next displays literal. Put

$$a_n = C_{B,W} \frac{1 + \log L}{L}, \qquad z_{p,q} = C_{B,W}\left(\frac{1}{Lp} + \frac{1}{Lq} + \frac{1}{n}\right),$$

where the constants are enlarged once below. Division by (7.44), together with (7.39) and (7.43), gives, whenever $rt_p + st_q \le 4$,

$$\begin{aligned}\left|\mathbb{P}_{C,S}(p^r q^s \mid m) - \frac{H(rt_p + st_q)}{p^r q^s}\right| \\ \ll_{B,W} \frac{a_n + z_{p,q}}{p^r q^s} + D_p(r)C_q(s) + p^{-r}D_q(s) + \frac{1}{n}.\end{aligned}$$

The one-prime restoration gives symmetrically

$$\left|\mathbb{P}_{C,S}(p^r \mid m) - \frac{H(rt_p)}{p^r}\right| \ll_{B,W} \frac{a_n + z_p}{p^r} + D_p(r) + \frac{1}{n}, \qquad z_p \ll_{B,W} \frac{1}{Lp} + \frac{1}{n}.$$

The required geometric sums are

$$\sum_{r\ge 2} \frac{1}{p^r} = \frac{1}{p(p-1)}, \qquad \sum_{r\ge 2} \frac{r}{p^r} = \frac{2p-1}{p(p-1)^2}, \qquad \sum_{r\ge 2} \frac{2r-3}{p^r} = \frac{p+1}{p(p-1)^2}; \tag{7.45}$$

in particular, all three are $O(p^{-2})$.

Combine (7.39) and (7.42)– (7.45), and sum these pointwise error majorants at fixed p, q before passing to a prime-row norm. There is a quantity

$$\alpha_{B,W}(n) \ll_{B,W} \frac{1 + \log L}{L}, \qquad \alpha_{B,W}(n) \log L \longrightarrow 0,$$

such that, throughout the indicated four-mark ranges,

$$\begin{aligned}\sum_{\substack{r\ge 2\\ rt_p + t_q \le 4}} \left|\mathbb{P}_{C,S}(p^r q \mid m) - \frac{H(rt_p + t_q)}{p^r q}\right| &\le \frac{\alpha_{B,W}(n)}{p^2 q} + O_{B,W}(L/n),\\ \sum_{\substack{s\ge 2\\ t_p + st_q \le 4}} \left|\mathbb{P}_{C,S}(pq^s \mid m) - \frac{H(t_p + st_q)}{pq^s}\right| &\le \frac{\alpha_{B,W}(n)}{pq^2} + O_{B,W}(L/n),\\ \sum_{\substack{r,s\ge 2\\ rt_p + st_q \le 4}} \left|\mathbb{P}_{C,S}(p^r q^s \mid m) - \frac{H(rt_p + st_q)}{p^r q^s}\right| &\le \frac{\alpha_{B,W}(n)}{p^2 q^2} + O_{B,W}(L^2/n).\end{aligned} \tag{7.46}$$

The one-prime restoration gives in exactly the same way

$$\begin{aligned}\sum_{\substack{r\ge 2\\ rt_p \le 4}} \left|\mathbb{P}_{C,S}(p^r \mid m) - \frac{H(rt_p)}{p^r}\right| &\le \frac{\alpha_{B,W}(n)}{p^2} + O_{B,W}(L/n),\\ \left|\mathbb{P}_{C,S}(q \mid m) - \frac{H(t_q)}{q}\right| &\le \frac{\alpha_{B,W}(n)}{q} + O_{B,W}(1/n).\end{aligned} \tag{7.47}$$

By symmetry, the high-power marginal needed in the IJ and JJ products is

$$\sum_{\substack{s\ge 2\\ st_q \le 4}} \left|\mathbb{P}_{C,S}(q^s \mid m) - \frac{H(st_q)}{q^s}\right| \le \frac{\alpha_{B,W}(n)}{q^2} + O_{B,W}(L/n); \tag{7.48}$$

the analogous first-power estimate at p follows by exchanging p, q in the second line of (7.47). All errors in these displays have been accounted for explicitly: (7.39) contributes $O_{B,W}((1+\log L)/L)$ times the reciprocal series, (7.42) contributes $O_{B,W}(1/L)$, and there are only $O_W(L)$, respectively $O_W(L^2)$, endpoint terms.

Let

$$\Gamma_{r,s} := \mathrm{Cov}_{C,S}(I_{p^r}, I_{q^s}), \qquad \mathcal{K}_{r,s} := \frac{H(rt_p + st_q) - H(rt_p)H(st_q)}{p^r q^s}.$$

Use $|ab - a_0b_0| \le |a - a_0||b| + |a_0||b - b_0|$ with (7.46)– (7.48). For the JJ marginal error only, enlarge the coupled chamber to the containing Cartesian product. The geometric sums in (7.45) then give the literal covariance error ledger

$$\sum_{\substack{r\ge 2\\ rt_p+t_q\le 4}} |\Gamma_{r,1} - \mathcal{K}_{r,1}| \le \frac{\alpha_{B,W}(n)}{p^2q} + O_{B,W}(L/n),$$

$$\sum_{\substack{s\ge 2\\ t_p+st_q\le 4}} |\Gamma_{1,s} - \mathcal{K}_{1,s}| \le \frac{\alpha_{B,W}(n)}{pq^2} + O_{B,W}(L/n),$$

$$\sum_{\substack{r,s\ge 2\\ rt_p+st_q\le 4}} |\Gamma_{r,s} - \mathcal{K}_{r,s}| \le \frac{\alpha_{B,W}(n)}{p^2q^2} + O_{B,W}(L^2/n).$$

Finally (7.41) and $\sum_{r\ge 2} r/p^r \ll p^{-2}$ show that the complete four-mark chamber contributes at most

$$C\frac{t_pt_q}{p^2q} + \frac{\alpha_{B,W}(n)}{p^2q} + O_{B,W}(L/n), \qquad C\frac{t_pt_q}{pq^2} + \frac{\alpha_{B,W}(n)}{pq^2} + O_{B,W}(L/n),$$
$$C\frac{t_pt_q}{p^2q^2} + \frac{\alpha_{B,W}(n)}{p^2q^2} + O_{B,W}(L^2/n), \tag{7.49}$$

in the JI, IJ, and JJ orientations. The constant C depends only on U, and in particular is independent of B, W.

Beyond-four-mark reciprocal tails. We now treat every forced product beyond y^4. The full tilted measure also has bounded density ratio, so

$$\mathbb{P}_{C,S}(D \mid m) \ll_{B,W} D^{-1} + n^{-1}. \tag{7.50}$$

For JI, if $rt_p + t_q > 4$, then

$$(r-2)t_p > 4 - t_q - 2t_p \ge 1,$$

so the reciprocal tail is at most $O_{B,W}(Ly^{-1}/(p^2q))$. The transpose is identical. For JJ, put $E_{r,s} = (r-2)t_p + (s-2)t_q$. The case $r = s = 2$ cannot occur. If exactly one exponent exceeds two, say $r \ge 3, s = 2$, then $rt_p > 2$ and

$$E_{r,2} = (r-2)t_p > \frac{2(r-2)}{r} \ge \frac{2}{3}.$$

If both exceed two, then $rt_p + st_q = 2(t_p + t_q) + E_{r,s} \le 3E_{r,s}$, so $E_{r,s} > 4/3$. The JJ reciprocal tail is consequently $O_{B,W}(L^2y^{-2/3}/(p^2q^2))$. The elementary inequality

$$|\,\mathrm{Cov}(\mathbf{1}_A, \mathbf{1}_B)| \le \mathbb{P}(A \cap B) + \mathbb{P}(A)\mathbb{P}(B)$$

and (7.50) also show that all literal endpoints are $O_{B,W}(L/n)$ in JI, IJ and $O_{B,W}(L^2/n)$ in JJ. Thus the complete tail and endpoint ledger is

	reciprocal tail	literal endpoint
JI	$O_{B,W}(Ly^{-1}/(p^2q))$	$O_{B,W}(L/n)$
IJ	$O_{B,W}(Ly^{-1}/(pq^2))$	$O_{B,W}(L/n)$
JJ	$O_{B,W}(L^2y^{-2/3}/(p^2q^2))$	$O_{B,W}(L^2/n)$.

(7.51)

Together with (7.49), this proves (7.34)–(7.35) with an intermediate pair error

$$\epsilon^{(0)}_{B,W}(n) \ll_{B,W} \frac{1+\log L}{L} + \frac{L}{y} + \frac{L^2}{y^{2/3}}, \qquad \epsilon^{(0)}_{B,W}(n)\log L \longrightarrow 0. \tag{7.52}$$

For the diagonal, repeat the omitted-score and exact-restoration argument with $E = \{p\}$. Since

$$J_p^2 = \sum_{r\ge 2}(2r-3)\mathbf{1}_{p^r|m},$$

(7.42) and the one-prime version of (7.39) give, for $rt_p \le 4$, an absolute leading bound C/p^2, an $\epsilon^{(0)}_{B,W}(n)/p^2$ error, and an $O_{B,W}(L^2/n)$ endpoint. If $rt_p > 4$, then $(r-2)t_p > 4 - 2t_p \ge 2$, so the remaining reciprocal tail is $O_{B,W}(L^2 y^{-2}/p^2)$. This proves (7.36) with a leading constant independent of B, W.

Aggregate row norm. Finally, sum the three off-diagonal bounds in the row norm. For fixed p, the contraction uses the following literal finite-sum factorizations (with $q = p$ omitted throughout):

$$p\sum_{q\ne p}\left(\frac{C_{\text{pow}}t_pt_q}{p^2q}, \frac{C_{\text{pow}}t_pt_q}{pq^2}, \frac{C_{\text{pow}}t_pt_q}{p^2q^2}\right) = C_{\text{pow}}\left(\frac{t_p}{p}\sum_{q\ne p}\frac{t_q}{q}, t_p\sum_{q\ne p}\frac{t_q}{q^2}, \frac{t_p}{p}\sum_{q\ne p}\frac{t_q}{q^2}\right).$$

The same identities with the t's removed factor the three error terms. Thus no interchange of an unscaled $o(1)$ with the prime sum is being used. These exact identities give the explicit orientation ledger

$$\mathcal{R}^{\text{main}}_{JI} \ll \frac{t_p}{p}\sum_q \frac{t_q}{q}, \qquad \mathcal{R}^{\text{err}}_{JI} \ll \frac{\epsilon^{(0)}_{B,W}(n)}{p}\sum_q\frac{1}{q},$$

$$\mathcal{R}^{\text{main}}_{IJ} \ll t_p\sum_q \frac{t_q}{q^2}, \qquad \mathcal{R}^{\text{err}}_{IJ} \ll \epsilon^{(0)}_{B,W}(n)\sum_q\frac{1}{q^2},$$

$$\mathcal{R}^{\text{main}}_{JJ} \ll \frac{t_p}{p}\sum_q \frac{t_q}{q^2}, \qquad \mathcal{R}^{\text{err}}_{JJ} \ll \frac{\epsilon^{(0)}_{B,W}(n)}{p}\sum_q\frac{1}{q^2}.$$

The leading terms use

$$\sum_{W<q\le y}\frac{t_q}{q} = O(1), \qquad \sum_{q>W}\frac{1}{q^2} \ll \frac{1}{W}.$$

The only harmonic loss is in the unweighted JI pair error, and (7.52) gives

$$\epsilon^{(0)}_{B,W}(n)\sum_{q\le y}\frac{1}{q} \ll \epsilon^{(0)}_{B,W}(n)\log L = o(1).$$

For the endpoints, $p \le y$, $\pi(y) \le y$, and the worst JJ term in (7.51) give explicitly

$$p\,\pi(y)\frac{L^2}{n} \le \frac{y^2L^2}{n} = o(1).$$

On the diagonal, $I_pJ_p = J_p$, $0 \le J_p \le J_p^2$, and hence

$$|\operatorname{Var}(V_p) - \operatorname{Var}(I_p)| \le 2|\operatorname{Cov}(I_p, J_p)| + \operatorname{Var}(J_p) \le 3\mathbb{E}J_p^2.$$

Thus its weighted row contribution is $O(1/p) \le O(1/W)$, with the same box-independent leading constant. With a sufficiently large constant $C_{B,W}$, define

$$\epsilon_{B,W}(n) := C_{B,W}\left\{\epsilon^{(0)}_{B,W}(n)\log L + \frac{y^2L^2}{n}\right\}.$$

This single function simultaneously dominates the pair and row remainders and tends to zero. This proves (7.37) and completes the proof. □

We retain the following direct row calculation as an explicit verification of the truncated marked range, every high-power tail, and every endpoint. Its crude B-dependent leading constants are not used to choose W; that box-uniform choice comes only from Lemma 7.5. The calculation also supplies the physical and coefficient-weighted corollaries used below.

First, (7.24) and Mertens' estimates give the J_pI_q orientation

$$p \sum_{\substack{W<q\le y\\ q\ne p}} \sum_{\substack{k\ge 2\\ p^k\le y^2}} |\operatorname{Cov}_{C,S}(I_{p^k}, I_q)|$$

$$\le C_B p \sum_{k\ge 2} \frac{1}{p^k} \sum_{q\le y} \frac{t_q}{q} + O_{B,W}\left\{ \frac{p}{L} \sum_{k\ge 2} \frac{1}{p^k} \sum_{q\le y} \frac{1}{q} \right\} \ll_B \frac{1}{p} + o_{W,n}(1/p). \tag{7.53}$$

For the transpose I_pJ_q,

$$p \sum_{\substack{W<q\le y\\ q\ne p}} \sum_{\substack{l\ge 2\\ q^l\le y^2}} |\operatorname{Cov}_{C,S}(I_p, I_{q^l})|$$

$$\le \left\{ C_B t_p + O_{B,W}(L^{-1}) \right\} \sum_{q>W} \sum_{l\ge 2} \frac{1}{q^l} \ll_B \frac{1}{W} + o_{W,n}(1). \tag{7.54}$$

Here $\sum t_q/q = O(1)$, $L^{-1}\sum 1/q = o(1)$, and $\sum_{q>W, l\ge 2} q^{-l} \ll W^{-1}$.

The bounded tilt has density ratio $O_B(1)$ relative to C, independently of W. Thus Proposition 7.1 gives, for distinct p, q and $p^kq^l \le y^4$,

$$|\operatorname{Cov}_{C,S}(I_{p^k}, I_{q^l})| \ll_B \frac{1}{p^kq^l} + O_{B,W}\left(\frac{1}{p^kq^lL}\right).$$

For arbitrary k, l, elementary common-multiple counting gives the fallback

$$|\operatorname{Cov}_{C,S}(I_{p^k}, I_{q^l})| \ll_{B,W} \frac{1}{p^kq^l} + \frac{1}{n}. \tag{7.55}$$

There are $O_W(L)$ relevant powers at each prime. Consequently the part with $p^kq^l \le y^4$ contributes

$$\ll_B \frac{1}{pW} + o_{W,n}(1).$$

For the reciprocal part beyond y^4, split according as $p^k > y^2$ or $p^k \le y^2$; in the latter case necessarily $q^l > y^2$. Hence

$$p \sum_{\substack{q>W,\ k,l\ge 2\\ p^kq^l>y^4}} \frac{1}{p^kq^l} \ll_W py^{-2} \sum_{q>W, l\ge 2} \frac{1}{q^l} + p \sum_{k\ge 2} \frac{1}{p^k} \sum_{\substack{q\le y,\ l\ge 2\\ q^l>y^2}} \frac{1}{q^l}$$

$$\ll_W \frac{1}{y}.$$

Using (7.55) on this part, the complete two-extra-power orientation is therefore

$$p \sum_{\substack{W<q\le y\\ q\ne p}} |\operatorname{Cov}_{C,S}(J_p, J_q)| \ll_B \frac{1}{pW} + o_{W,n}(1) + O_{B,W}\left(\frac{pyL^2}{n}\right).$$

The endpoint term is uniform and is $o_n(1)$, since $p \le y = n^{2/9}$.

The portions omitted from (7.53) and (7.54) are also explicit. From (7.55),

$$p \sum_{q\le y} \sum_{p^k>y^2} |\operatorname{Cov}_{C,S}(I_{p^k}, I_q)| \ll_{B,W} py^{-2} \sum_{q\le y} \frac{1}{q} + O_{B,W}\left(\frac{pyL}{n}\right) = o_n(1),$$

$$p\sum_{q\le y}\sum_{q^l>y^2}|\operatorname{Cov}_{C,S}(I_p,I_{q^l})|\ll_{B,W} y^{-1}+O_{B,W}\left(\frac{pyL}{n}\right)=o_n(1).$$

Thus every endpoint contribution has the advertised form $O(pyL^{O(1)}/n)=o(1)$.

Finally consider the same-prime diagonal. Since

$$J_p^2=\sum_{k\ge 2}(2k-3)I_{p^k},$$

the marked bound through $p^k\le y^4$, followed by the elementary tail bound, gives

$$\mathbb{E}_{C,S}J_p^2\ll_B \frac{1}{p^2}+O_{B,W}\left(\frac{1}{p^2L}+\frac{L}{y^4}+\frac{L^2}{n}\right).$$

Moreover $I_pJ_p=J_p$, so

$$\begin{aligned}p\,|\mathrm{Var}_{C,S}(V_p)-\mathrm{Var}_{C,S}(I_p)|&\le p\{2|\operatorname{Cov}_{C,S}(I_p,J_p)|+\mathrm{Var}_{C,S}(J_p)\}\\&\ll_B \frac{1}{p}+o_{W,n}(1)+O_{B,W}(pL^2/n).\end{aligned}$$

The preceding displays control every high-power and endpoint tail. The box-uniform leading terms are supplied by Lemma 7.5; together they give

$$\|\operatorname{Cov}_{C,S}(V_p,V_q)-\operatorname{Cov}_{C,S}(I_p,I_q)\|_{\mathrm{row}}\le\frac{C_{\mathrm{pow}}}{W}+\epsilon_{B,W}(n).\tag{7.56}$$

To record explicitly the relative estimate used later, suppose that a coefficient vector satisfies

$$\|b\|_\infty+\sum_p\frac{|b_p|}{p}\le C_{\mathrm{cmp}}w,\qquad \sum_p\frac{b_p^2}{p}\le C_{\mathrm{cmp}}w^2,\tag{7.57}$$

where C_{cmp} is fixed independently of the compact tilt box and before W is enlarged. Put

$$A_1:=\sum_p\frac{t_p|b_p|}{p},\quad A_2:=\sum_p\frac{t_p|b_p|}{p^2},\quad B_1:=\sum_p\frac{|b_p|}{p},\quad B_2:=\sum_p\frac{|b_p|}{p^2}.$$

Since $0<t_p\le 1$ and $p>W$,

$$A_1,B_1\le C_{\mathrm{cmp}}w,\qquad A_2,B_2\le\frac{C_{\mathrm{cmp}}}{W}w.\tag{7.58}$$

The product-weighted bounds (7.34)–(7.35) now give, with their reciprocal tails already included in $\epsilon_{B,W}(n)$,

$$\begin{aligned}&\sum_{p\ne q}|b_pb_q|\,\{|\operatorname{Cov}(J_p,I_q)|+|\operatorname{Cov}(I_p,J_q)|\}\\&\quad\le C_{\mathrm{pow}}(A_2A_1+A_1A_2)+\epsilon_{B,W}(n)(B_2B_1+B_1B_2)\\&\quad\le\frac{C(C_{\mathrm{cmp}})}{W}w^2+C(C_{\mathrm{cmp}})\epsilon_{B,W}(n)w^2.\end{aligned}$$

Similarly, (7.35) and (7.36) contribute

$$C_{\mathrm{pow}}A_2^2+\epsilon_{B,W}(n)B_2^2\quad\text{and}\quad\{C_{\mathrm{pow}}+\epsilon_{B,W}(n)\}\sum_p\frac{b_p^2}{p^2},$$

respectively. By (7.57) and (7.58), these are $O_{C_{\mathrm{cmp}}}(w^2/W)$ plus $C(C_{\mathrm{cmp}})\epsilon_{B,W}(n)w^2$. Keeping the harmless numerical factors visible, and enlarging C_{pow} once, the entire non-endpoint contribution is bounded by

$$\frac{3C_{\mathrm{pow}}(C_{\mathrm{cmp}}^2 + C_{\mathrm{cmp}})}{W} w^2 + 3(C_{\mathrm{cmp}}^2 + C_{\mathrm{cmp}})\epsilon_{B,W}(n)w^2.$$

Indeed the off-diagonal products use $A_2, B_2 \le C_{\mathrm{cmp}} w/W$, while $\sum_p b_p^2/p^2 \le W^{-1} \sum_p b_p^2/p$; the last comparison uses only $W^{-1} \le 1$ and the nonnegativity of $C_{\mathrm{cmp}}^2 + C_{\mathrm{cmp}}$. This is the scalar contraction recorded in the finite Lean audit, rather than an implicit nonlinear use of the O-notation.

It remains only to sum the literal endpoints, rather than fold them into an unscaled $o(1)$. The coefficient bound gives

$$\sum_p |b_p| \le \pi(y)\|b\|_\infty \le C_{\mathrm{cmp}} wy.$$

Consequently the JI/IJ, JJ, and diagonal endpoints are, respectively,

$$\frac{L}{n} \sum_{p \ne q} |b_p b_q| \ll C_{\mathrm{cmp}}^2 w^2 \frac{y^2 L}{n},$$
$$\frac{L^2}{n} \sum_{p \ne q} |b_p b_q| \ll C_{\mathrm{cmp}}^2 w^2 \frac{y^2 L^2}{n},$$
$$\frac{L^2}{n} \sum_p b_p^2 \ll C_{\mathrm{cmp}}^2 w^2 \frac{y L^2}{n}.$$

All three are $o_n(w^2)$. Define, for this fixed comparison constant,

$$\epsilon^{\mathrm{rel}}_{B,W,C_{\mathrm{cmp}}}(n) := C(C_{\mathrm{cmp}}) \left\{ \epsilon_{B,W}(n) + \frac{y^2 L^2}{n} \right\}. \tag{7.59}$$

After increasing the fixed factor $C(C_{\mathrm{cmp}})$, this also contains the smaller y^2L/n and yL^2/n endpoint terms. Hence $\epsilon^{\mathrm{rel}}_{B,W,C_{\mathrm{cmp}}}(n) \to 0$ for every fixed B, W, C_{cmp}. The complete reciprocal tails and all coefficient factors are now recorded in this single relative remainder. Consequently

$$\left| \mathrm{Var}_{C,S}\left(\sum_p b_p V_p \right) - \mathrm{Var}_{C,S}\left(\sum_p b_p I_p \right) \right| \le \frac{C_{\mathrm{rel}}(C_{\mathrm{cmp}})}{W} w^2 + \epsilon^{\mathrm{rel}}_{B,W,C_{\mathrm{cmp}}}(n) w^2, \tag{7.60}$$

where $C_{\mathrm{rel}}(C_{\mathrm{cmp}})$ is independent of the tilt-box radius. Formula (7.59) records explicitly the fixed coefficient loss and all endpoint and local-fugacity remainders; the threshold for its convergence may depend on B, W, C_{cmp}. Thus no unrecorded transpose or diagonal term is being absorbed in the relative estimate used in Section 8.

The Stieltjes form of (7.18), followed by the same local expansion, gives for $p^k \le y^2$

$$|\operatorname{Cov}_{C,S}(\mathbf{1}_{p^k|m}, \phi(m/n))| \ll_{B,W} \frac{\|\phi\|_{BV}}{p^k L}.$$

For the omitted exponents, comparison of the tilted and untilted densities and the elementary count of multiples give, uniformly for $W < p \le y$,

$$\sum_{p^k > y^2} \mathbb{P}_{C,S}(p^k \mid m) \ll_{B,W} y^{-2} + \frac{L}{n} = o\left(\frac{1}{pL}\right).$$

Indeed, after division by the cell size, the geometric multiple-counting term is $O(y^{-2})$, while the $+1$ endpoint is summed over at most $O(L)$ exponents. Relative to $1/(pL)$, the two displayed terms have ratios at most L/y and yL^2/n, respectively, and both tend to zero. Since the covariance of a divisibility indicator with the bounded physical test is at most twice its probability times $\|\phi\|_\infty$, summation over every exponent now proves

$$|\operatorname{Cov}_{C,S}(v_p(m), \phi(m/n))| \ll_{B,W} \frac{\|\phi\|_{BV}}{pL}. \tag{7.61}$$

A further physical factor $e^{\zeta\phi(m/n)/L} = 1 + O_B(1/L)$ is handled by the same expansion. Reweighting the finitely many exact head cells in a compact positive box also preserves every estimate, because their normalized main marked profile in (7.3) is common. We have therefore proved uniform one- and two-prime kernels, the physical Stieltjes column, and the aggregate $O(1/W)$ prime-power transfer throughout the complete compact band, physical, and finite-head homotopy used in the next section. At no point was an arbitrary residue class or a $U = 5/2$ terminal estimate invoked.

For later reference we isolate the preceding stability statement, including the effect of the guarded deletions.

Corollary 7.6 (Compact full tilt and guarded cells)**.** *Let $\nu = \sum_{i=1}^{m} \lambda_i \nu_i$ be a fixed finite cell mixture as in (7.1), with its nonzero weights in a fixed compact subset of the relative interior. Reweight it by*

$$\frac{d\nu_\xi}{d\nu}(m) = \frac{1}{Z_\xi} \exp\left\{ \frac{1}{L} \left(S(m) + \zeta\phi(m/n) + \sum_{\ell=1}^{e} \kappa_\ell H_\ell^\circ(m) \right) \right\}, \tag{7.62}$$

where the effective prime fugacities, ζ, and the finitely many κ_ℓ lie in a fixed box, ϕ is bounded with bounded variation, and the centered head indicators are bounded. Then

(i) *there are constants $0 < c_{B,W} < C_{B,W} < \infty$ such that $c_{B,W} \le d\nu_\xi/d\nu \le C_{B,W}$ pointwise;*

(ii) *all squarefree, prime-power, physical, and coefficient-weighted estimates in this section hold uniformly under ν_ξ, with the same box-independent coefficients of $1/W$;*

(iii) *if a set $\mathcal{R}_n^{\mathrm{guard}}$ of at most $O(yL^A)$ numerical guards is deleted before normalization, the change in every normalized one- or two-mark prime-row operator, and hence in its harmonically averaged band operator, is*

$$O_{B,W}\left(\frac{y^2 L^{A+3}}{n} \right) = o_n(\alpha_0), \qquad \alpha_0 \asymp \frac{1}{\log L}. \tag{7.63}$$

In particular the deletion error remains relative at the moving low cell, not merely $o_n(1)$.

Proof. The elementary inequality

$$\Omega(m) \log W \le \log(C_* n)$$

gives $|S(m)|/L \le B/\log W + O_{B,W}(1/L)$. The physical and head scores in (7.62) are $O_{B,W}(1/L)$. Normalization therefore proves (i). Taylor expansion of the latter two factors has a uniform $O_{B,W}(1/L)$ remainder. The Stieltjes formula makes their leading physical average cancel in centered marked moments, while the common marked profile of the head cells makes the between-cell term cancel. Combining this with Lemma 7.2 and Lemma 7.5, together with (7.60) proves (ii). We shall also use the following component-level consequence of the same calculation. Let $\nu_{i,\xi}$ be the conditional law on the i-th exact component cell before the numerical guards are deleted. Then, uniformly over the finitely many nonzero components and the fixed box,

$$\sup_{i,i'} \left| \mathbb{E}_{\nu_{i,\xi}} V_p - \mathbb{E}_{\nu_{i',\xi}} V_p \right| \ll_{B,W} \frac{1}{pL} \qquad (W < p \le y). \tag{7.64}$$

Indeed (7.11) proves this before tilting. For $p^k \le y^2$, the one-local specialization of the omitted-local-score expansion and exact fugacity restoration used in Lemma 7.5 writes every component mean as the same restored local model, with error

$$O_{B,W}\left(\frac{1}{p^k L} + \frac{1}{n}\right).$$

Summing k gives $O_{B,W}(1/(pL) + L/n) = O_{B,W}(1/(pL))$, since $pL^2/n \le yL^2/n = o(1)$. The omitted powers are $o(1/(pL))$ by the tail calculation immediately preceding (7.61). A physical tilt is $1 + O_B(1/L)$, and a head tilt is constant on each exact component, so neither changes this comparison. Thus (7.64) is a displayed consequence of (ii), rather than an implicit appeal to a common profile.

We first make the component-weight bookkeeping literal. If the later construction prescribes weights λ_i for the individually guarded component laws, put

$$s_i := \nu_i((\mathcal{R}_n^{\mathrm{guard}})^c), \qquad \widetilde{\lambda}_i := \frac{\lambda_i/s_i}{\sum_j \lambda_j/s_j}.$$

Every component has $\Theta_W(n)$ candidates, so

$$1 - s_i \ll_W \frac{yL^A}{n} = o(1)$$

uniformly in the fixed finite component family. Hence all s_i are positive and the weights $\widetilde{\lambda}_i$ remain in a fixed compact subset of the relative interior whenever the prescribed λ_i's do. The exact measure identity

$$\left(\sum_i \widetilde{\lambda}_i \nu_i\right)\left(\cdot \mid (\mathcal{R}_n^{\mathrm{guard}})^c\right) = \sum_i \lambda_i \nu_i \left(\cdot \mid (\mathcal{R}_n^{\mathrm{guard}})^c\right)$$

follows because the conditioned weight of component i on the left is $\widetilde{\lambda}_i s_i / \sum_j \widetilde{\lambda}_j s_j = \lambda_i$. The identity remains exact after applying and normalizing the same positive exponential tilt on both sides. We may therefore apply the pre-deletion estimates to the compact mixture with weights $\widetilde{\lambda}_i$, and recover precisely the prescribed guarded mixture after conditioning. This also covers the simpler case in which the guards are deleted directly from a globally weighted mixture.

Let ν_ξ denote the resulting pre-deletion tilted law, $\widehat{\nu}_\xi$ the actual guarded tilted law, and put

$$\delta_\xi := \nu_\xi(\mathcal{R}_n^{\mathrm{guard}}).$$

Every component cell has $\Theta_W(n)$ candidates, and (i) makes the probability of a single candidate $O_{B,W}(1/n)$. Hence, uniformly on the fixed box,

$$\delta_\xi \ll_{B,W} \frac{yL^A}{n} = o(1). \tag{7.65}$$

We use the following elementary conditional-deletion identities. If μ is a probability law, G has $\mu(G) = \delta < 1$, and $\widehat{\mu} = \mu(\cdot \mid G^c)$, then, for every bounded F,

$$\mathbb{E}_{\widehat{\mu}} F - \mathbb{E}_\mu F = \frac{\delta \mathbb{E}_\mu F - \mathbb{E}_\mu(F \mathbf{1}_G)}{1-\delta}, \qquad \left|\mathbb{E}_{\widehat{\mu}} F - \mathbb{E}_\mu F\right| \le \frac{2\delta}{1-\delta}\|F\|_\infty. \tag{7.66}$$

Writing the covariance difference as

$$\begin{aligned}\mathrm{Cov}_{\widehat{\mu}}(F,H) - \mathrm{Cov}_\mu(F,H) = \{&\mathbb{E}_{\widehat{\mu}}(FH) - \mathbb{E}_\mu(FH)\} \\ &- \{\mathbb{E}_{\widehat{\mu}} F - \mathbb{E}_\mu F\}\mathbb{E}_{\widehat{\mu}} H\end{aligned}$$

$$- \mathbb{E}_\mu F\{\mathbb{E}_{\widehat{\mu}} H - \mathbb{E}_\mu H\}$$

and applying (7.66) three times gives the explicit recentered bound

$$\left|\mathrm{Cov}_{\widehat{\mu}}(F,H) - \mathrm{Cov}_\mu(F,H)\right| \le \frac{6\delta}{1-\delta}\|F\|_\infty\|H\|_\infty. \tag{7.67}$$

Thus both normalization and the two changes of means are included.

Apply this with $\mu = \nu_\xi$, $G = \mathcal{R}_n^{\mathrm{guard}}$, and $V_p = v_p(m)$. On every candidate, for all sufficiently large n,

$$pV_p(m) \ll yL, \qquad \sum_{W<r\le y} V_r(m) \le \frac{\log m}{\log W} \ll_W L. \tag{7.68}$$

Let

$$E_{pr}^{\mathrm{guard}} := \mathrm{Cov}_{\widehat{\nu_\xi}}(V_p, V_r) - \mathrm{Cov}_{\nu_\xi}(V_p, V_r).$$

For fixed p, duality with $H_b = \sum_r b_r V_r$, $\|b\|_\infty \le 1$, followed by (7.67)– (7.68), gives

$$\begin{aligned} p\sum_{W<r\le y} |E_{pr}^{\mathrm{guard}}| &= \sup_{\|b\|_\infty\le 1} p\left|\mathrm{Cov}_{\widehat{\nu_\xi}}(V_p, H_b) - \mathrm{Cov}_{\nu_\xi}(V_p, H_b)\right| \\ &\ll_{B,W} \delta_\xi y L^2 \ll_{B,W} \frac{y^2L^{A+2}}{n} \le C_{B,W}\frac{y^2L^{A+3}}{n}. \end{aligned} \tag{7.69}$$

The expectation bound in (7.66) proves the analogous one-mark estimate; replacing either V by $I_p = \mathbf{1}_{p|m}$, or the second score by a bounded physical or finite-head score, only makes the envelope smaller.

It remains to identify the norm after band aggregation. For $\|b\|_\infty \le 1$, the covariance-error contribution on band i is bounded exactly by

$$\begin{aligned} \frac{1}{H_i}\left|\sum_{p\in\mathcal{P}_i}\sum_r E_{pr}^{\mathrm{guard}} b_{j(r)}\right| &\le \frac{1}{H_i}\sum_{p\in\mathcal{P}_i}\frac{1}{p}\left(p\sum_r |E_{pr}^{\mathrm{guard}}|\right) \\ &\le \sup_u u\sum_r |E_{ur}^{\mathrm{guard}}|. \end{aligned}$$

Denote this raw normalized band covariance-error operator by $E_{\mathrm{band}}^{\mathrm{guard}}$. Thus harmonic averaging is a contraction of the normalized prime-row norm (7.33); the uniformly bounded arithmetic gauge projections do not alter its order. For completeness, the sharp conjugation used later also follows from the displayed relative rate. Put $S_n = \mathrm{diag}(\alpha_0,\ldots,\alpha_s)$. If $\|q\|_\infty \le 1$, apply the raw row estimate to $b_j = \alpha_j q_j$, for which $\|b\|_\infty \le 1$. Dividing output row i by α_i gives

$$\left|(S_n^{-1}E_{\mathrm{band}}^{\mathrm{guard}} S_n q)_i\right| \ll_{B,W} \frac{y^2L^{A+3}}{n\alpha_i} \le C_{B,W}\frac{y^2L^{A+3}}{n\alpha_0} = o(1).$$

Indeed, once the cells are nonempty, $\alpha_0 < \delta < \alpha_j$ for every $j \ge 1$: the low cell is contained in $[t_0, \delta]$, every positive cell lies in $(\delta, 1]$, and nonemptiness follows for the fixed finite permitted mesh from the prime number theorem. Thus α_0 is the smallest cell center. The arithmetic quantities H_i, α_i are prime geometry and hence are unchanged by deleting candidates. The gauge projections are uniformly bounded in this sharp norm by (8.41). This proves (7.63). Finally $y^2L^{A+3}/n = n^{-5/9}L^{A+3} = o(1/\log L) = o(\alpha_0)$. □

8 The smooth-row bridge and exact finite-band fitting

In this section we complete the fractional construction left by Section 6. We use throughout

$$L = \log n, \qquad N = \frac{n}{L}, \qquad \theta = \frac{2}{9}, \qquad y = n^\theta, \qquad U = \theta^{-1} = \frac{9}{2}.$$

The real cutoff W, the head set $F_{\mathrm{hd}} = \{p : p \le W\}$, and the head modulus P_{hd} have already been fixed. All constants in this section may depend on W, but not on n or on a prime in $(W, y]$.

8.1 The integer row ledger and the active physical target

Recall that $Y = T/D'$ is the target remaining after the anchor divisor has been removed, and that B_0 is the combined product of the fixed lower factors and the base states of the precharged switch banks. The rough-stage residual is

$$r_p = v_p(Y) - v_p(B_0) - \sum_{a\in\mathcal{A}} x_a v_p(a)$$

as in (6.51). In the trivial complete rough-signature row the fractional selector has the additive decomposition

$$x^{\mathrm{sm}} = x^{\mathrm{top}} + f^{\mathrm{prot}} + z^{\mathrm{act}}. \tag{8.1}$$

Here x^{top} is the constant-density component on the short top interval $H = (2n - Kh, 2n]$, f^{prot} is the broad layer of weight β_{prot}/L reserved for the tangent construction, and z^{act} is the adjustable broad excess. The support of z^{act} is contained in

$$J = (n, 2n - Kh],$$

and is therefore disjoint from the support of x^{top}. The protected summand may share numerical coordinates with z^{act}; in that case the final inclusion weight is their sum. Only z^{act} is differentiated below. In particular, neither x^{top} nor f^{prot} is ever folded into the bridge probability measure.

We first record the integer-versus-real distinction which will also be needed in the final rounding. Freeze B_0, all nonsmooth fractional rows, the whole top component, the protected floor, and every other smooth seed not used in the interpolation below. Let $m_{\mathrm{sm,fix}}$ be the number of fixed factors of B_0 in the smooth signature row, and let $m_{\mathrm{sm,fr}}$ be the total mass of all frozen contributions in that row. Put

$$m_{\mathrm{sm,oth}} := m_{\mathrm{sm,fr}} - m_{\mathrm{sm,fix}} - \sum_a x_a^{\mathrm{top}} - \sum_a f_a^{\mathrm{prot}}.$$

Let $\widetilde{q}_0 = \sum_m z_m^{\mathrm{act}}$ be the actual active mass left by the guarded rough construction, before any interpolation in this section. Thus the actual post-guard mass in the smooth row is $m_{\mathrm{sm,fr}} + \widetilde{q}_0$. Define, without an asymptotic replacement,

$$Q_{\mathrm{sm}}(0) := \operatorname*{argmin}_{k\in\mathbb{Z}} |k - (m_{\mathrm{sm,fr}} + \widetilde{q}_0)|, \qquad q_0 := Q_{\mathrm{sm}}(0) - m_{\mathrm{sm,fr}};$$

break a possible tie by taking the smaller integer. In particular, $|q_0 - \widetilde{q}_0| \le 1/2$. Choose the integer total smooth-row quota $Q_{\mathrm{sm}}(d) = Q_{\mathrm{sm}}(0) - d$, and define

$$\begin{aligned} q_{\mathrm{sm}}^{\mathrm{flex}}(d) &:= Q_{\mathrm{sm}}(d) - m_{\mathrm{sm,fix}} \in \mathbb{Z}, \\ q^{\mathrm{act}}(d) &:= Q_{\mathrm{sm}}(d) - m_{\mathrm{sm,fr}} = q_0 - d. \end{aligned} \tag{8.2}$$

Thus $q^{\mathrm{act}}(d)$ is a real mass; it need not be an integer. The exact ledger identity is

$$\boxed{Q_{\mathrm{sm}}(d) = m_{\mathrm{sm,fix}} + \sum_a x_a^{\mathrm{top}} + \sum_a f_a^{\mathrm{prot}} + q^{\mathrm{act}}(d) + m_{\mathrm{sm,oth}}, \qquad q_{\mathrm{sm}}^{\mathrm{flex}}(d) \in \mathbb{Z}.} \tag{8.3}$$

Equivalently, the total weight of the flexible numerical coordinates in the smooth row is $q_{\mathrm{sm}}^{\mathrm{flex}}(d)$. There is no separate integrality requirement on any of the four real summands on the right-hand side of (8.3).

We now define the two frozen ledgers without identifying fractional mass with a cardinality. Let $\mathscr{F}_0$ be the set of frozen integral factors (so that its contribution includes B_0), and let $\mathscr{X}_0$ be the set of all frozen fractional coordinates, with their actual weights denoted by x_a^0. The active coordinates z^{act} are excluded from $\mathscr{X}_0$. Put, exactly,

$$m_0 := \#\mathscr{F}_0 + \sum_{a\in\mathscr{X}_0} x_a^0, \qquad \Lambda_0 := \sum_{a\in\mathscr{F}_0} \log a + \sum_{a\in\mathscr{X}_0} x_a^0 \log a.$$

Thus m_0 is the exact *total frozen mass*, generally not an integer, whereas Λ_0 is its ordinary logarithmic contribution. In particular $\log B_0$ is included in Λ_0, while $\log D'$ is not, since D' has already been divided from the target. The source of the mass estimate is explicit. The charged fixed tokens together with the flexible mass in every nontrivial row total exactly t_R. Summing those identities and adjoining the actual post-guard smooth raw mass gives

$$m_0 + \widetilde{q}_0 = \sum_{R \neq 1} t_R + \{\text{actual post-guard smooth raw mass}\} = h + O(N/L).$$

The nearest-integer initialization changes this by only $O(1)$. Together with the corresponding height estimates from the rough stage, this gives

$$m_0 + q_0 = h + O(N/L), \qquad \Lambda_0 = m_0 L + O(N), \qquad \log Y = hL + O(N).$$

The first relation uses the same exact convention for every retained factor, bank base state, guard, and fractional row. Consequently

$$A_0 := \log Y - \Lambda_0 - q_0 L = O(N).$$

For $q^{\text{act}}(d) = q_0 - d$, the height above n which the active smooth component must supply is

$$A(d) := \log Y - \Lambda_0 - q^{\text{act}}(d)L = A_0 + dL. \tag{8.4}$$

Fix once and for all a number $\mu \in (0, \log 2)$, and let d be a nearest integer to

$$d_* := \frac{\mu q_0 - A_0}{L + \mu}. \tag{8.5}$$

The bounds above give $d = O(N/L)$. Since $q_0 \asymp N$, they also give

$$q^{\text{act}}(d) \asymp N, \qquad \frac{A(d)}{q^{\text{act}}(d)} = \mu + O(L/N).$$

This centers the mean of the *remaining active smooth mass*; it is not a statement about the mean of all lower factors. In particular, retained exceptional factors and bank states have already been charged in Λ_0.

Here is the promised literal feasibility of the two mass changes. Before freezing the head cells, keep two disjoint positive-length subintervals of the zero-head stratum

$$\{m \in J : P^+(m) \leq y,\ (m, P_{\text{hd}}) = 1\},$$

one below and one above $e^{\mu} n$. The rough active density is a fixed positive multiple of $1/L$ there, and the unused ceiling is bounded below by a fixed constant. Each subpool consequently has both removal capacity and addition capacity $\gg_W N$. First change their combined mass by $q_0 - \widetilde{q}_0 = O(1)$, and then by $-d = O(N/L)$, splitting each change between the two pools. Spreading these changes uniformly alters an individual active coordinate by $O_W(L^{-2})$. All head valuations remain zero, no coordinate leaves $[0, 1]$, and the two physical splitting parameters remain a fixed distance from their endpoints. Thus the mass after this operation is exactly $q^{\text{act}}(d) = q_0 - d$. The later two-pool interpolation changes the split, but not this mass. This proves, rather than assumes, that every integer d used in (8.5) is attainable.

8.2 Exact head cells and physical interpolation

After the choice of d, subtract from the target each frozen head valuation and let A_p^{act} be the valuation still assigned to the active smooth component at $p \leq W$. This quantity has a useful coordinatewise verification. Put

$$\gamma_{p,W} := \begin{cases} (c - C_0)/(3(p-1)), & p \in F_{\text{anc}}, \\ c/(2(p-1)), & p \in F_{\text{hd}} \setminus F_{\text{anc}}. \end{cases}$$

The nonsmooth flexible point, the protected layer, and the initial active smooth point are all F_{hd}-free. Hence the only head charge besides D' is the already combined base B_0. From (4.23), its preserved form (5.16), (6.47), and (6.48), we have simultaneously

$$\begin{aligned} A_p^{\mathrm{act}} &= v_p(T) - v_p(D') - v_p(B_0) \\ &\geq \left\{ \gamma_{p,W} - C_W \frac{\delta_*}{\theta p} \right\} N - o_W(N) \geq 2\epsilon_W N \qquad (p \leq W), \end{aligned} \tag{8.6}$$

after the prescribed choice of δ_* and then sufficiently large n. Also $A_p^{\mathrm{act}} \leq v_p(T) \leq C_W N$. Changing d uses only the zero-head pool just displayed, so it changes every one of these targets by zero (and, even without that choice, an $O_W(N/L)$ change would be smaller than the fixed linear margin). Consequently

$$\epsilon_W N \leq A_p^{\mathrm{act}} \leq C_W N \qquad (p \leq W).$$

Choose a fixed integer E_W so large that

$$\frac{(A_p^{\mathrm{act}})_{p \leq W}}{q^{\mathrm{act}}(d)} \in \operatorname{int} \operatorname{conv}\{0, E_W \mathbf{e}_p : p \leq W\} \tag{8.7}$$

with a fixed positive barycentric margin. The margin follows directly from (8.6), not from a limiting interiority assertion. Explicitly, the coefficient of $E_W \mathbf{e}_p$ is $A_p^{\mathrm{act}}/(E_W q^{\mathrm{act}}(d))$, and the coefficient of 0 is one minus their sum.

The vertices in (8.7) are realized by exact valuation/coprimality patterns: the zero vertex is $v_\ell(m) = 0$ for every $\ell \leq W$, and the p-vertex is

$$v_p(m) = E_W, \qquad v_\ell(m) = 0 \quad (\ell \leq W,\ \ell \neq p).$$

These are finite divisibility inclusion–exclusion cells, not residue classes. By Proposition 7.1, each such pattern has $\Theta_W(n)$ y-smooth members in every fixed positive-length subinterval of $(n, 2n]$.

Choose $\eta_0 > 0$ so that

$$0 < \mu - \eta_0 < \mu + \eta_0 < \log 2$$

and then fix constants

$$1 < a_- < b_- < e^{\mu - \eta_0}, \qquad e^{\mu + \eta_0} < a_+ < b_+ < 2.$$

For every n, define the scaled intervals

$$I_-(n) := (a_- n, b_- n], \qquad I_+(n) := (a_+ n, b_+ n].$$

Since $(2 - b_+)n > Kh$ for all sufficiently large n, one has $I_+(n) \subset (n, 2n - Kh]$; the corresponding containment for $I_-(n)$ is immediate. In each head pattern reserve positive-density smooth subpools in both $I_-(n)$ and $I_+(n)$. Assign the barycentric masses from (8.7), subtracting exactly the mass assigned from the active zero-pattern stratum. Thus no new $\Theta(N)$ mass is added. Split each pattern mass between its two physical subpools. The two logarithmic means lie on opposite sides of μ; hence a compact choice of the splitting coefficients makes the total active logarithm equal to (8.4) exactly. The $O(L)$ error caused by rounding d_* is removed by changing one splitting coefficient by $O(L/N) = o(1)$.

Every assigned cell has $\Theta_W(n)$ candidates and receives $O_W(N)$ mass. Its individual weights are therefore $O_W(1/L)$. The barycentric and physical splitting margins leave a positive-density zero-pattern active component with weights bounded below by a fixed multiple of $1/L$. This component, rather than the frozen protected floor, supplies all covariance lower bounds below.

The Stieltjes estimate (7.61) also shows that the physical and head redistribution changes a marked valuation at $p > W$ by at most

$$O_W\left(\frac{N}{pL}\right). \tag{8.8}$$

We summarize this construction as a measure-realization lemma, so the probability space used by the bridge is not left implicit.

Lemma 8.1 (Baseline active-measure realization)**.** *Let $\mathscr{E} = \{0\} \cup \{E_W \mathbf{e}_p : p \le W\}$ be the finite head-pattern simplex above, and let $\sigma \in \{-,+\}$ indicate the physical interval $I_\sigma(n)$. Put $q_n := q^{\mathrm{act}}(d)$. After deleting all numerical guards, there are disjoint structured cells $C_{e,\sigma,n}$ and masses $q_{e,\sigma,n} > 0$ such that*

$$\nu_n = \sum_{e\in\mathscr{E}} \sum_{\sigma\in\{-,+\}} \lambda_{e,\sigma}(n)\nu_{e,\sigma,n}, \qquad \lambda_{e,\sigma}(n) := \frac{q_{e,\sigma,n}}{q_n}, \tag{8.9}$$

is a probability measure, where $\nu_{e,\sigma,n}$ is uniform on $C_{e,\sigma,n}$. More precisely:

(i) *the vector $(\lambda_{e,\sigma}(n))$ stays in a fixed compact subset of the open simplex;*

(ii) *$\#C_{e,\sigma,n} = \Theta_W(n)$, also after guard deletion;*

(iii) *the coordinate realization*

$$z_m^0 = \frac{q_{e,\sigma,n}}{\#C_{e,\sigma,n}} \quad (m \in C_{e,\sigma,n}) \tag{8.10}$$

satisfies $z_m^0 = O_W(1/L)$, and on a fixed positive fraction of the zero-head cells one has $z_m^0 \ge c_W/L$;

(iv) *after addition of the frozen protected weight at a shared clean coordinate, the total weight is still $O_W(1/L) < 1$ for large n; and*

(v) *the total mass is exactly q_n, all head moments are exactly $(A_p^{\mathrm{act}})_{p\le W}$, and the ordinary logarithmic moment is exactly $A(d)$.*

Proof. For a fixed head pattern and sign, let $C_{e,\sigma,n}$ be the smooth members of the corresponding scaled interval with that exact valuation and coprimality pattern, after disjointification and guard deletion. The marked cell estimate gives $\Theta_W(n)$ members before deletion. By Lemma 6.5, the whole excluded set has size $O_c(N) + O(yL) = o(n)$ (and its intersection with the active smooth row is smaller). Hence the same $\Theta_W(n)$ estimate remains after deletion. The strict barycentric margin in (8.7) and the strict two-sided physical split put every normalized assigned mass in a fixed compact subset of the open simplex. This proves (i), (ii), and (8.9).

Since $q_n \asymp n/L$, formula (8.10) gives the upper bound in (iii); the compact lower mass and the upper candidate-count bound give the asserted zero-head lower bound. These cells lie below $2n - Kh$, hence are disjoint from the top component, and all other unavailable coordinates were deleted as guards. The only allowed overlap is with the protected layer, whose weight is itself $O_W(1/L)$. This proves (iv). Finally the barycentric equations are exactly the head equations, the two physical split parameters were chosen to solve the ordinary logarithm, and their weights sum to q_n. These are finite linear identities, proving (v). $\square$

8.3 The Poisson–Dickman covariance operator

We next prove the analytic inverse used for the band fit. Let Π be the scale-invariant Poisson process on $(0,1]$, conditioned on $\sum_{t\in\Pi} t = U = 9/2$. Set

$$h(s) = \frac{\rho(U-s)}{\rho(U)}, \qquad K(s,t) = \frac{\rho(U-s-t)}{\rho(U)} - h(s)h(t),$$

with $\rho(v) = 0$ for $v < 0$. The Palm identities (2.15)–(2.16) imply that the covariance of the additive statistic $\sum_{t\in\Pi} f(t)$ is the quadratic form of

$$(Af)(s) = h(s)f(s) + \int_0^1 K(s,t)f(t)\,\frac{dt}{t}.$$

Put $F(x) = \rho(U-x)/\rho(U)$. The Dickman delay equation shows that $F \in C^2[0,2]$: indeed $U - x \in [5/2, 9/2]$, a compact interval on which the two differentiated delay equations are continuous. We have $h(s) = F(s)$ and

$$K(s,t) = F(s+t) - F(s)F(t).$$

Since $F(0) = 1$, a mixed first difference followed by Taylor's theorem gives

$$|K(s,t)| \ll st \qquad (0 \le s,t \le 1). \tag{8.11}$$

Lemma 8.2 (Poisson–Dickman quotient inverse)**.** *On $\mathcal{H} = L^2((0,1], dt/t)$, the operator A is nonnegative, Fredholm, and*

$$\ker A = \operatorname{span}\{t\}.$$

Consequently there is a constant $\kappa_0 > 0$ such that

$$\langle f, Af\rangle_{\mathcal{H}} \ge \kappa_0 \inf_{\lambda\in\mathbb{R}} \int_0^1 |f(t) - \lambda t|^2\,\frac{dt}{t}. \tag{8.12}$$

On bounded functions one also has

$$\inf_{\lambda\in\mathbb{R}} \|f - \lambda t\|_\infty \ll \|Af\|_\infty. \tag{8.13}$$

Moreover, on

$$X_t := \{f : f(t)/t \in L^\infty(0,1]\}, \qquad \|f\|_{X_t} := \|f(t)/t\|_\infty,$$

one has the weighted quotient estimate

$$\inf_{\lambda\in\mathbb{R}} \|f - \lambda t\|_{X_t} \ll \left\|\frac{Af}{t}\right\|_\infty. \tag{8.14}$$

Proof. Let

$$\mathcal{D} := \{f \in L^\infty(0,1] : f = 0 \text{ a.e. on } (0,\varepsilon) \text{ for some } \varepsilon > 0\}.$$

This is dense in $\mathcal{H} = L^2((0,1], dt/t)$. For $f \in \mathcal{D}$, the additive sum $S_f = \sum_{x\in\Pi} f(x)$ has only finitely many nonzero terms. The one- and two-point Palm formulas therefore give, first for $f, g \in \mathcal{D}$,

$$\begin{aligned}\operatorname{Cov}_U(S_f, S_g) &= \int_0^1 h(t)f(t)g(t)\,\frac{dt}{t} \\ &\quad + \int_0^1\int_0^1 K(s,t)f(s)g(t)\,\frac{ds}{s}\frac{dt}{t} = \langle f, Ag\rangle_{\mathcal{H}}.\end{aligned} \tag{8.15}$$

Multiplication by h is boundedly invertible on $\mathcal{H}$. Also, by (8.11),

$$\int_0^1\int_0^1 |K(s,t)|^2\,\frac{ds}{s}\frac{dt}{t} < \infty.$$

Thus the integral part is Hilbert–Schmidt, the bilinear form (8.15) extends continuously from $\mathcal{D}$ to $\mathcal{H}$, and A is the resulting bounded self-adjoint operator. Nonnegativity on the dense core passes to the extension. Since A is an invertible multiplication operator plus a compact operator, it is Fredholm of index zero.

If $Af = 0$, the integral equation and (8.11) give the explicit Cauchy–Schwarz bound

$$|f(s)| \le \frac{1}{\min h} \left(\int_0^1 |K(s,t)|^2 \frac{dt}{t} \right)^{1/2} \|f\|_{\mathcal{H}} \ll s \|f\|_{\mathcal{H}}$$

for almost every s. Here $\min_{[0,1]} h > 0$, since $U - s \in [7/2, 9/2]$. Change the representative on a null set so that this holds for every $s \in (0,1]$. Then

$$\sum_{x \in \Pi} |f(x)| \ll \sum_{x \in \Pi} x = U \qquad \mathbb{P}_U\text{-almost surely.}$$

The same bound makes all one- and two-point second-moment integrals finite. Truncating at $x \ge \varepsilon$ and using dominated convergence extends (8.15) to this particular f. Consequently

$$\mathrm{Var}_U(S_f) = \langle f, Af \rangle_{\mathcal{H}} = 0.$$

Thus S_f is a genuine almost-sure constant: for some c,

$$Z_c(\pi) := \sum_{x \in \pi} f(x) - c = 0 \quad \text{for } \mathbb{P}_U\text{-almost every } \pi.$$

We spell out the Palm disintegration which turns this almost-sure statement into an additive identity. Campbell–Mecke, disintegrated with respect to the total mass, says for every nonnegative measurable Φ

$$\mathbb{E}_U \sum_{r \in \Pi} \Phi(r, \Pi \setminus \{r\}) = \frac{1}{\rho(U)} \int_0^1 \rho(U - r) \mathbb{E}_{U-r} \Phi(r, \Pi) \, \frac{dr}{r}, \tag{8.16}$$

$$\mathbb{E}_U \sum_{\substack{r,s \in \Pi \\ r \ne s}} \Phi(r, s, \Pi \setminus \{r, s\}) = \frac{1}{\rho(U)} \int_0^1 \int_0^1 \rho(U - r - s) \mathbb{E}_{U-r-s} \Phi(r, s, \Pi) \frac{dr}{r} \frac{ds}{s}. \tag{8.17}$$

Terms with a negative residual mass are interpreted as zero. These formulas follow first with all atoms at least ε, where the sums are finite, and then by monotone convergence. Apply (8.17) to $(Z_c(\pi) + f(r) + f(s))^2$ and (8.16) to $(Z_c(\pi) + f(u))^2$. Their left sides vanish. Since $\rho(U - r - s) > 0$ for $r + s \le 1$, we obtain, for almost every such pair (r, s), under the *same* residual law $\mathbb{P}_{U-r-s}$,

$$\begin{aligned} Z_c(\Pi) + f(r) + f(s) &= 0, \\ Z_c(\Pi) + f(r + s) &= 0. \end{aligned}$$

For the second equality we used the one-point formula at $u = r + s$; the pullback of its null set under $(r, s) \mapsto r + s$ is null by Fubini. Subtraction gives $f(r) + f(s) = f(r + s)$ for almost every $r, s > 0$ with $r + s \le 1$. For completeness, the measurable local Cauchy step can be seen by convolution. On each compact interval $I \Subset (0,1)$, choose $\psi \in C_c^\infty(0, 1 - \sup I)$ with $\int \psi = 1$, and put

$$F_0(x) := \int \{f(x + y) - f(y)\} \psi(y) \, dy.$$

The bound $f(t) = O(t)$ gives local integrability, so translation continuity in L^1_{loc} makes F_0 continuous. Fubini and the almost-everywhere Cauchy identity give $F_0 = f$ almost everywhere. If two compact intervals overlap, their continuous representatives agree almost everywhere on the overlap and hence, by continuity, agree everywhere there. They therefore glue to a continuous representative on $(0,1)$. The identity $F_0(x+z) = F_0(x) + F_0(z)$ holds first almost everywhere and then everywhere on each compact subtriangle by continuity. A continuous local additive function is λx; the glued representative shows that the same λ applies on all overlapping intervals. Thus $f(t) = \lambda t$ almost everywhere. Conversely this function is in the kernel because the conditioned sum is constant. Since zero is now an isolated eigenvalue of a Fredholm self-adjoint operator, (8.12) follows.

For completeness, put

$$\widetilde{K}(s,t) := \frac{K(s,t)}{t},$$

using at $t = 0$ the continuous value $F'(s) - F(s)F'(0)$. On $L^\infty(0,1]$,

$$Af = M_h f + \mathcal{K}_\infty f, \qquad (\mathcal{K}_\infty f)(s) = \int_0^1 \widetilde{K}(s,t) f(t)\, dt.$$

The image of the L^∞ unit ball under $\mathcal{K}_\infty$ is uniformly bounded and equicontinuous in $C[0,1]$. Arzelà–Ascoli therefore makes $\mathcal{K}_\infty$ compact as an operator into L^∞. Since M_h is invertible, A is Fredholm on L^∞. If a bounded f lies in its kernel, the integral equation first gives $f(t) = O(t)$, so the preceding L^2 kernel calculation applies and

$$\ker_{L^\infty} A = \operatorname{span}\{t\}.$$

The induced map

$$\overline{A} : L^\infty / \operatorname{span}\{t\} \longrightarrow \operatorname{Ran} A, \qquad [f] \longmapsto Af,$$

is a bounded bijection between Banach spaces, because a Fredholm range is closed. The quotient norm here is $\|[f]\| = \inf_{\lambda\in\mathbb{R}} \|f - \lambda t\|_\infty$. The open mapping theorem gives (8.13).

For the weighted assertion write $f(t) = tQ(t)$. Then

$$\frac{A(tQ)(s)}{s} = h(s)Q(s) + \int_0^1 \frac{K(s,t)}{s} Q(t)\, dt. \tag{8.18}$$

The quotient kernel in (8.18) is even more explicit:

$$\frac{K(s,t)}{s} \longrightarrow F'(t) - F'(0)F(t) \quad (s \downarrow 0), \tag{8.19}$$

uniformly in $t \in [0,1]$. Since $F \in C^2[0,2]$, this defines a continuous kernel on the whole compact square. The corresponding conjugated operator

$$(\mathcal{B}Q)(s) := \frac{A(tQ)(s)}{s} = h(s)Q(s) + \int_0^1 \frac{K(s,t)}{s} Q(t)\, dt$$

maps the L^∞ unit ball in its integral part into a uniformly bounded, equicontinuous subset of $C[0,1]$, and that part is compact by Arzelà–Ascoli. Thus $\mathcal{B}$ is an invertible multiplication operator plus a compact operator. Its kernel consists of constants, because $Q \mapsto tQ$ is an isometry from L^∞ onto X_t and $tQ \in \ker A$ exactly when Q is constant. The induced map from $L^\infty/\mathbb{R}$, equipped with the quotient norm $\inf_{\lambda\in\mathbb{R}} \|Q - \lambda\|_\infty$, to the closed range of $\mathcal{B}$ is a bounded bijection, and the open mapping theorem gives (8.14). □

8.4 Finite bands, the low row, and arithmetic transfer

Put

$$t_p := \frac{\log p}{\log y}, \qquad t_0 := \frac{\log W}{\log y}.$$

Choose fixed $0 < \delta < 1$ and a fixed regular relative partition

$$[t_0, \delta] = I_0, \qquad (\delta, 1] = I_1 \sqcup \cdots \sqcup I_s.$$

Thus, after merging a possible terminal fragment with its neighbor,

$$c_{\mathrm{mesh}} \eta \inf I_j \le |I_j| \le \eta \inf I_j \qquad (j \ge 1) \tag{8.20}$$

for fixed $c_{\mathrm{mesh}} > 0$ and $\eta > 0$. In the exponent convention used in the tangent section, the low endpoint is $\delta_b = \theta\delta$, and

$$\theta(\delta + \eta) = \delta_b + \theta\eta.$$

We write

$$\begin{aligned} \mathcal{P}_j &:= \{p : W < p \le y,\ t_p \in I_j\}, \\ \Omega_j(m) &:= \sum_{p \in \mathcal{P}_j} v_p(m), \\ H_j &:= \sum_{p \in \mathcal{P}_j} \frac{1}{p}, \qquad \alpha_j := \frac{1}{H_j} \sum_{p \in \mathcal{P}_j} \frac{t_p}{p}. \end{aligned}$$

Let $D = \mathrm{diag}(H_0, \ldots, H_s)$.

The active point constructed in Section 8.2 is supported on a fixed finite disjoint union $\mathcal{S}_n$ of positive-length physical intervals with exact finite head patterns; any initial overlap is disjointified before normalization. Denote its weights by z_m^0, and normalize

$$q_n = \sum_{m \in \mathcal{S}_n} z_m^0 = q^{\mathrm{act}}(d) \asymp N, \qquad \nu_n(m) := \frac{z_m^0}{q_n}.$$

Every z_m^0 is $O_W(1/L)$, and on a fixed positive-density zero-pattern subpool it is bounded below by c_W/L. By Lemma 6.5, only the promoted smooth anchors and bank endpoints can meet these active smooth cells, so the relevant numerical guards number $O(yL)$. By Corollary 7.6, deleting them changes every normalized marked row by

$$O\left(\frac{y^2 L^{O(1)}}{n}\right) = o(\alpha_0).$$

This uses the actual guard geometry, rather than an assertion about an arbitrary $o(n)$ set.

For a compact vector of parameters, reweight only the active mass by

$$z_m(\xi) = q_n \frac{z_m^0 \exp\Big\{L^{-1}\Big(\sum_j \eta_j \Omega_j(m) + \zeta R(m) + \sum_\nu \kappa_\nu H_\nu^\circ(m)\Big)\Big\}}{\displaystyle\sum_{\ell \in \mathcal{S}_n} z_\ell^0 \exp\left\{L^{-1}\left(\sum_j \eta_j \Omega_j(\ell) + \zeta R(\ell) + \sum_\nu \kappa_\nu H_\nu^\circ(\ell)\right)\right\}}, \tag{8.21}$$

where $R(m) = \log(m/n)$. To define the head scores, choose one reference pattern $e_* \in \mathscr{E}$, put

$$H_e(m) := \mathbf{1}_{C_{e,-,n} \cup C_{e,+,n}}(m), \qquad H_e^\circ(m) := H_e(m) - \mathbb{E}_{\nu_n} H_e \quad (e \ne e_*),$$

and enumerate these $|\mathscr{E}| - 1$ centered indicators as $H_1^\circ, \ldots, H_e^\circ$. Changing the centering adds only a constant to the exponential score and hence does not change the tilted measure. The denominator keeps the active mass exactly q_n. For any two statistics F, G,

$$\frac{\partial}{\partial \xi_G} \sum_m z_m(\xi) F(m) = \frac{q_n}{L} \mathrm{Cov}_{\nu_{n,\xi}}(F, G).$$

Here and below a *compact effective-tilt box* means that, after use of the exact logarithmic identity, every coefficient of $v_p(m)$ for $p > W$, and every remaining head or physical coefficient, has absolute value at most a fixed B. This is the hypothesis of (7.13). It does not impose a bounded value on an unidentified redundant parameter. In particular the whole low-band contribution to the exponent is bounded by

$$\frac{B}{L} \sum_{p > W} v_p(m) \le \frac{B}{\log W} + O_B(L^{-1}),$$

so the fact that $H_0 \sim \log L$ creates no divergent count tilt.

We first state the quotient inverse, including the moving low row.

For a partition $\mathcal{P}$, define the step injection and harmonic averaging operators on $\mathbb{R}^{s+1}$ by

$$(J_{\mathcal{P}}b)(t) = b_j \quad (t \in I_j), \qquad (E_{\mathcal{P}}f)_j = \frac{1}{H_j^{(c)}} \int_{I_j} f(t)\,\frac{dt}{t},$$

$$B_{t_0,\mathcal{P}} := E_{\mathcal{P}} A^{(t_0)} J_{\mathcal{P}}, \qquad A^{(t_0)} f(s) := h(s) f(s) + \int_{t_0}^{1} K(s,t) f(t)\,\frac{dt}{t}.$$

Here $H_j^{(c)} = \int_{I_j} dt/t$, $D_c = \operatorname{diag}(H_j^{(c)})$, and $\alpha_j^{(c)} = E_{\mathcal{P}}(t)_j$. Thus

$$H_0^{(c)} = \log(\delta/t_0), \qquad \alpha_0^{(c)} = \frac{\delta - t_0}{H_0^{(c)}}.$$

The continuum projection and gauge are

$$P_{\alpha,c} b = b - \alpha^{(c)} \frac{(\alpha^{(c)})^T D_c b}{(\alpha^{(c)})^T D_c \alpha^{(c)}}, \qquad \mathcal{G}_c := \{b : (\alpha^{(c)})^T D_c b = 0\}. \tag{8.22}$$

The arithmetic projection and gauge, which are different finite-n objects, are

$$P_{\alpha,n} b = b - \alpha \frac{\alpha^T D b}{\alpha^T D \alpha}, \qquad \mathcal{G}_n := \{b : \alpha^T D b = 0\}.$$

From this point on $\mathcal{G}$ without a subscript means $\mathcal{G}_n$; the continuum gauge is always written $\mathcal{G}_c$. Because the two displayed denominators are bounded below, both projections are uniformly bounded in ℓ^∞. We prove below, for each fixed permitted mesh, that $\|P_{\alpha,n} - P_{\alpha,c}\|_{\infty\to\infty} = o_n(1)$.

Let $Z = (R, H_1^\circ, \dots, H_e^\circ)$, with one head indicator deleted so that the constant score is absent, and put

$$\Gamma_{\xi,n} := \operatorname{Cov}_{\nu_{n,\xi}}(Z, Z).$$

When n is understood we write $\Gamma_\xi = \Gamma_{\xi,n}$; in particular $\Gamma_{0,n} := \operatorname{Cov}_{\nu_n}(Z,Z)$ is the actual finite-n baseline covariance, not a putative limiting covariance. Define the genuine head/physical Schur band matrix

$$\mathcal{C}_\xi^Z := q_n \left\{ \operatorname{Cov}_\xi(\Omega,\Omega) - \operatorname{Cov}_\xi(\Omega, Z)\Gamma_\xi^{-1} \operatorname{Cov}_\xi(Z,\Omega) \right\}. \tag{8.23}$$

There is no "row Schur complement": normalization in (8.21) has already removed the constant row, whose covariance would be zero.

We shall use the following finite-dimensional perturbation fact twice, once for the arithmetic transfer and once for the nuisance Schur block.

Lemma 8.3 (Stable projected Schur perturbation)**.** *Let D be positive diagonal, let P be the D-orthogonal projection onto a gauge space $\mathcal{G}$, and let B be symmetric. Assume*

$$b^T B b \ge \kappa b^T D b \quad (b \in \mathcal{G}), \qquad \left\| \left(P D^{-1} B P|_{\mathcal{G}} \right)^{-1} \right\|_{\infty\to\infty} \le C.$$

If a symmetric perturbation E satisfies

$$|b^T E b| \le \varepsilon b^T D b \quad (b \in \mathcal{G}), \qquad \|P D^{-1} E P\|_{\infty\to\infty} \le \varepsilon, \tag{8.24}$$

then, whenever $\varepsilon < \min(\kappa/2, (2C)^{-1})$, $B - E$ has quadratic gap at least $\kappa/2$ on $\mathcal{G}$, and

$$\left\| \left(P D^{-1} (B - E) P|_{\mathcal{G}} \right)^{-1} \right\|_{\infty\to\infty} \le 2C.$$

Proof. The quadratic assertion follows by subtracting the first bound in (8.24). On $\mathcal{G}$, factor the projected operator as

$$PD^{-1}(B-E)P = PD^{-1}BP\left\{I-(PD^{-1}BP|_{\mathcal{G}})^{-1}PD^{-1}EP\right\}.$$

The second factor differs from the identity by an operator of ℓ^∞-norm at most $C\varepsilon < 1/2$, so its Neumann series has norm at most two. This proves the inverse assertion. ☐

Uniformity convention for the next two lemmas. The continuum constants used below depend only on $U = 9/2$, the regularity constant c_{mesh}, and a fixed upper bound for $w := \delta+\eta$; in particular they are chosen before W, a tilt box, or a particular mesh is instantiated. The constants multiplying the nonvanishing arithmetic errors W^{-1} are those already obtained in the marked estimates and are independent of every later tilt-box radius and mesh. We first choose W large enough to absorb those terms and then choose w below the continuum smallness threshold, as required by the global order in (6.2). After a radius B and one permitted finite mesh $\mathcal{P}$ have been fixed, the notation

$$r_{n;B,W,\mathcal{P}} = o_{n;B,W,\mathcal{P}}(1)$$

means convergence uniformly for ξ in that B-box and uniformly over the unit ball of the vector norm occurring in the estimate. For a quadratic-form estimate the unit ball means $b^TDb \le 1$. Its threshold may depend on $(B, W, \delta, \eta, \mathcal{P})$; no uniform threshold over unbounded boxes or meshes varying with n is used. Finally, the continuum comparison object is itself $B_{t_0,\mathcal{P}}$ with the same $t_0 = (\log W)/(\log y)$ as the arithmetic operator at that n. Thus the argument compares two moving finite-dimensional objects at each n, rather than claiming convergence to a single fixed band matrix.

Lemma 8.4 (Moving-low-cell arithmetic quotient)**.** *There are structural thresholds $W_0 < \infty$ and $w_0 > 0$. Choose $W \ge W_0$ and then fixed $\delta, \eta > 0$ with $w := \delta+\eta \le w_0$. There are constants $\kappa > 0$ and $C_W < \infty$, independent of the permitted regular mesh and of the radius B of a subsequently fixed compact effective-tilt box, with the following quantifiers: for every fixed B and every fixed permitted mesh $\mathcal{P}$, there is $n_0 = n_0(B, W, \delta, \eta, \mathcal{P})$ such that the conclusions below hold for all $n \ge n_0$ and every ξ in the B-box. No common threshold over arbitrarily fine meshes or over unbounded boxes is asserted. Then Γ_ξ is positive definite and*

$$b^T\mathcal{C}_\xi^Z b \ge \kappa q_n \inf_{\lambda\in\mathbb{R}} \sum_j H_j|b_j-\lambda\alpha_j|^2, \tag{8.25}$$

$$\left\|\left(P_{\alpha,n}D^{-1}\frac{\mathcal{C}_\xi^Z}{q_n}P_{\alpha,n}\,|_{\alpha^TDb=0}\right)^{-1}\right\|_{\ell^\infty\to\ell^\infty} \ll_W 1. \tag{8.26}$$

The constants κ, C_W are uniform as $t_0 \downarrow 0$ and over all regular meshes satisfying (8.20); only the threshold n_0 may depend on the particular fixed mesh and box.

Proof. Step 1: the continuum inverse and form gap, uniformly in the moving low cell. This step contains no primes and no tilt parameter. In particular its constants are independent of W, B, n, and of the locations of the positive-cell endpoints subject to (8.20). Write $f = J_{\mathcal{P}}b$. Since $h(s) = 1+O(s)$ and $|K(s,t)| \le Cst$, direct harmonic averaging on $I_0 = [t_0, \delta]$ gives

$$(B_{t_0,\mathcal{P}}b)_0 = \{1+O(\delta)\}b_0 + O\left(\alpha_0^{(c)}\right)\|b\|_\infty; \tag{8.27}$$

indeed

$$\frac{1}{H_0^{(c)}}\int_{I_0} s\,\frac{ds}{s} = \alpha_0^{(c)}, \qquad \int_{t_0}^1 t|f(t)|\,\frac{dt}{t} \le \|b\|_\infty.$$

The C^2-regularity of F also gives

$$|h'(s)| \le C, \qquad |\partial_s K(s,t)| \le Ct.$$

Hence, on a positive relative cell I_j, the oscillation of $A^{(t_0)}f$ is at most $C\eta\|b\|_\infty$. These are ordinary supremum-norm estimates and their constants do not involve the growing harmonic mass $H_0^{(c)}$.

Suppose the projected inverse were not uniform. There would be $t_{0,\nu} \downarrow 0$, permitted partitions, and gauge-fixed vectors b_ν such that

$$\|b_\nu\|_\infty = 1, \qquad \|P_{\alpha,c} B_{t_{0,\nu},\mathcal{P}_\nu} b_\nu\|_\infty = o(1).$$

By (8.22), for bounded scalars $\beta_\nu = O(1)$,

$$B_{t_{0,\nu},\mathcal{P}_\nu} b_\nu = \beta_\nu \alpha^{(c)} + o_{\ell^\infty}(1). \tag{8.28}$$

More explicitly, applying $I - P_{\alpha,c}$ gives

$$\beta_\nu = \frac{(\alpha^{(c)})^T D_c B_{t_{0,\nu},\mathcal{P}_\nu} b_\nu}{(\alpha^{(c)})^T D_c \alpha^{(c)}} + o(1).$$

The denominator is bounded below by the cells meeting $[1/2, 1]$, while $\sum_j H_j^{(c)} \alpha_j^{(c)} = 1 - t_0$ and the compressed operators are uniformly bounded. Hence $\beta_\nu = O(1)$. Formula (8.27) now gives

$$|b_{\nu,0}| = O(\delta) + o(1). \tag{8.29}$$

On each positive cell, the oscillation estimate turns (8.28) into

$$A^{(t_{0,\nu})} J_{\mathcal{P}_\nu} b_\nu(s) = \beta_\nu s + O(\delta + \eta) + o(1).$$

On the low cell both sides are $O(\delta) + o(1)$, by (8.29) and (8.11), so the same estimate holds there. Extend the low value to $(0, t_{0,\nu})$ by $b_{\nu,0} t / t_{0,\nu}$. For $s \ge t_{0,\nu}$, the newly inserted integral term is bounded by

$$\int_0^{t_{0,\nu}} |K(s,t)|\, |b_{\nu,0}| \frac{t}{t_{0,\nu}} \frac{dt}{t} \ll |b_{\nu,0}| s t_{0,\nu} = O(t_{0,\nu}).$$

For $0 < s < t_{0,\nu}$, the diagonal term is $O(|b_{\nu,0}|) = O(\delta) + o(1)$, while the integral term is $O(s)$; this is already within the error allowed below. Thus, for the resulting bounded function $\widetilde{f}_\nu$,

$$A\widetilde{f}_\nu = \beta_\nu t + O(\delta + \eta) + o(1) \quad \text{in } L^\infty. \tag{8.30}$$

Pairing with the kernel vector t shows $\beta_\nu = O(\delta + \eta) + o(1)$. Then (8.13) puts $\widetilde{f}_\nu$ within $O(\delta + \eta) + o(1)$ of $\lambda_\nu t$. Finally, $(\alpha^{(c)})^T D_c b_\nu = 0$ is precisely $\int_{t_0}^1 t J_{\mathcal{P}} b_\nu \, dt/t = 0$; the linear extension changes this pairing by $O(t_0)$. Hence $\lambda_\nu = O(\delta + \eta) + o(1)$, contradicting $\|b_\nu\|_\infty = 1$ once $\delta + \eta$ is below a fixed threshold. For the quadratic estimate normalize $\|b\|_{D_c} = 1$ and use the same linear extension on $(0, t_0)$. Its extra squared norm is

$$\int_0^{t_0} \left| b_0 \frac{t}{t_0} \right|^2 \frac{dt}{t} = \frac{|b_0|^2}{2} \le \frac{1}{2H_0^{(c)}} = o(1),$$

and its change in the t-pairing is at most $|b_0| t_0/2 = o(1)$. If the compressed quadratic gap failed, the quadratic forms, and not only the norms, compare as follows. With $f = J_{\mathcal{P}} b$ on $[t_0, 1]$, harmonic averaging gives the exact identity

$$b^T D_c B_{t_0,\mathcal{P}} b = \int_{t_0}^1 f(s) (A^{(t_0)} f)(s) \, \frac{ds}{s}. \tag{8.31}$$

If $e(s) = b_0 s/t_0$ on $(0, t_0)$, then $|K(s,t)| \ll st$, Cauchy–Schwarz, and the preceding norm bound give

$$\begin{aligned}
\Big|\int_0^{t_0} h(s)e(s)^2\,\frac{ds}{s}\Big| &\ll |b_0|^2 = o(1),\\
\Big|\int_0^{t_0}\int_{t_0}^{1} K(s,t)e(s)f(t)\,\frac{ds}{s}\frac{dt}{t}\Big| &\ll |b_0|t_0\|f\|_{\mathcal{H}} = o(1),\\
\Big|\int_0^{t_0}\int_0^{t_0} K(s,t)e(s)e(t)\,\frac{ds}{s}\frac{dt}{t}\Big| &\ll |b_0|^2 t_0^2 = o(1).
\end{aligned}\tag{8.32}$$

Consequently

$$\langle \widetilde{f}, A\widetilde{f}\rangle_{\mathcal{H}} - b^T D_c B_{t_0,\mathcal{P}} b = o(1) \tag{8.33}$$

under unit D_c-normalization. If the compressed quadratic gap failed, the extended functions would therefore have $L^2(dt/t)$-norm $1 + o(1)$, asymptotically zero pairing with t, and $\langle \widetilde{f}, A\widetilde{f}\rangle = o(1)$. The Poisson–Dickman Poincaré inequality (8.12) would put them $o(1)$-close to $\operatorname{span}\{t\}$, while the gauge would force the scalar to be $o(1)$, a contradiction. This proves the continuum quadratic gap uniformly as the low endpoint moves to zero. We have therefore proved, with structural constants $C_c < \infty$ and $\kappa_c > 0$,

$$\begin{aligned}
\Big\|\Big(P_{\alpha,c} B_{t_0,\mathcal{P}} P_{\alpha,c}\,|_{\mathcal{G}_c}\Big)^{-1}\Big\|_{\infty\to\infty} &\le C_c,\\
b^T D_c B_{t_0,\mathcal{P}} b \ge \kappa_c\, b^T D_c b \qquad &(b \in \mathcal{G}_c).
\end{aligned}\tag{8.34}$$

Both constants are uniform for all sufficiently small $t_0 > 0$, hence uniform as $t_0 \downarrow 0$, and for every permitted regular mesh once $w \le w_0$. This is precisely the range used because $t_0 = (\log W)/(\log y) \to 0$. These are the continuum results; no arithmetic object has yet been used.

Step 2: arithmetic cell data and normalized operator convergence. We now give the discrete transfer rather than hiding it in a convergence phrase. Put $I_p = 1_{p|m}$ and $\Omega_i^{\mathrm{sf}} = \sum_{p\in\mathcal{P}_i} I_p$. For a prime coefficient vector $a = (a_p)$, define the squarefree and full normalized prime operators and band averaging by

$$\begin{aligned}
(\mathsf{A}^{\mathrm{sf}}_{n,\xi} a)_p &:= p\,\mathrm{Cov}_{\nu_{n,\xi}}\Big(I_p, \sum_q a_q I_q\Big),\\
(\mathsf{A}_{n,\xi} a)_p &:= p\,\mathrm{Cov}_{\nu_{n,\xi}}\Big(v_p, \sum_q a_q v_q\Big), \qquad\qquad (\mathsf{E}_n a)_i := \frac{1}{H_i}\sum_{p\in\mathcal{P}_i}\frac{a_p}{p}.
\end{aligned}$$

Thus

$$D^{-1}\,\mathrm{Cov}_\xi(\Omega^{\mathrm{sf}}, \Omega^{\mathrm{sf}}) = \mathsf{E}_n \mathsf{A}^{\mathrm{sf}}_{n,\xi} J_n,$$

where $(J_n b)_p = b_{j(p)}$. The squarefree marked formulas give a diagonal multiplier plus a compact kernel and an error E_{pq}:

$$\begin{aligned}
\mathrm{Var}(1_{p|m}) &= \frac{h(t_p)}{p} + E_{pp},\\
\mathrm{Cov}(1_{p|m}, 1_{q|m}) &= \frac{K(t_p,t_q)}{pq} + E_{pq} \quad (p \ne q), &(8.35)\\
\sup_{W<p\le y}\sum_{W<q\le y} p|E_{pq}| &\le \frac{C_{\mathrm{sf}}}{W} + \epsilon^{\mathrm{sf}}_{B,W}(n). &(8.36)
\end{aligned}$$

Here C_{sf} is absolute and $\epsilon^{\mathrm{sf}}_{B,W}(n) \to 0$ for each fixed compact tilt box. The $1/W$ term contains the squarefree diagonal correction $-h(t_p)^2/p^2$; equivalently one may include it in the exact diagonal multiplier. The leading functions are the same h, K throughout a compact effective-tilt box: Lemma 7.2 puts the entire tilt dependence in E_{pq}. Thus the continuum kernel is independent of the limiting fugacity. The definition of the box above and $\sum_{p>W} v_p/L \le 1/\log W + O(1/L)$ are exactly what make the constants in this assertion uniform.

The passage from the prime row norm to the band norm is the exact calculation

$$\left| \frac{1}{H_i} \sum_{p\in\mathcal{P}_i} \frac{1}{p} \sum_q pE_{pq} b_{j(q)} \right| \le \frac{1}{H_i} \sum_{p\in\mathcal{P}_i} \frac{1}{p} \sum_q p|E_{pq}|\,|b_{j(q)}|. \tag{8.37}$$

Thus the large value $H_0 \asymp \log L$ cancels rather than multiplying the error. Mertens' theorem gives, uniformly cell by cell,

$$H_0 = \log L + O_{W,\delta}(1), \qquad \alpha_0 = \frac{\delta + o(1)}{H_0}.$$

We record precisely which discrete cell data converge to their continuum counterparts. On every positive cell, quantitative Mertens summation gives

$$H_j - H_j^{(c)} = o_n(1), \qquad H_j\alpha_j - H_j^{(c)}\alpha_j^{(c)} = o_n(1) \quad (j \ge 1). \tag{8.38}$$

On the moving low cell the first absolute convergence need not hold and is not used: the fixed lower cutoff can leave a bounded Mertens constant. The correct statements are

$$H_0 = H_0^{(c)} + O_W(1), \quad \frac{H_0}{H_0^{(c)}} = 1 + o_n(1), \quad H_0\alpha_0 = \delta - t_0 + o_n(1) = H_0^{(c)}\alpha_0^{(c)} + o_n(1). \tag{8.39}$$

Indeed the last relation is partial summation applied to $(\log y)^{-1} \sum_{W<p\le y^\delta} (\log p)/p$. It follows, for the fixed finite mesh, that

$$\begin{aligned} \max_j |\alpha_j - \alpha_j^{(c)}| &= o_n(1), & \max_j \left| \frac{\alpha_j}{\alpha_j^{(c)}} - 1 \right| &= o_n(1), \\ \alpha^T D\alpha - (\alpha^{(c)})^T D_c \alpha^{(c)} &= o_n(1), & \|P_{\alpha,n} - P_{\alpha,c}\|_{\infty\to\infty} &= o_n(1). \end{aligned} \tag{8.40}$$

For the relative assertion, use (8.39) on the low cell; on every positive cell, $\alpha_j^{(c)} \gg \delta$, so the absolute convergence is already relative. This is exactly the additional information needed in the sharp norm. Put

$$S_n := \operatorname{diag}(\alpha_0, \ldots, \alpha_s), \qquad S_c := \operatorname{diag}(\alpha_0^{(c)}, \ldots, \alpha_s^{(c)}).$$

Then

$$\|S_n S_c^{-1} - I\|_{\infty\to\infty} = o_n(1), \qquad \|S_n^{-1} P_{\alpha,n} S_n - S_c^{-1} P_{\alpha,c} S_c\|_{\infty\to\infty} = o_n(1). \tag{8.41}$$

Indeed, if $\omega_{n,j} = H_j\alpha_j^2$ and $\omega_{c,j} = H_j^{(c)}(\alpha_j^{(c)})^2$, then

$$S_n^{-1} P_{\alpha,n} S_n = I - \mathbf{1}\frac{\omega_n^T}{\mathbf{1}^T\omega_n}, \qquad S_c^{-1} P_{\alpha,c} S_c = I - \mathbf{1}\frac{\omega_c^T}{\mathbf{1}^T\omega_c}.$$

The positive-cell coordinates of $\omega_n - \omega_c$ tend to zero by (8.38) and the relative center convergence. On the low cell, both coordinates are $(\delta - t_0)^2/H_0^{(c)} + o(1) = o(1)$. Their total masses converge to the same positive limit by (8.40), proving the second operator convergence. Thus the continuum and arithmetic gauges are compared only through quantities which genuinely converge; no assertion $\|D - D_c\| \to 0$ is made at the moving low cell.

For clarity, here is the double prime-band quadrature as a normalized operator calculation. Put

$$(TJ_{\mathcal{P}}b)(s) := \int_{t_0}^{1} K(s,t)(J_{\mathcal{P}}b)(t)\,\frac{dt}{t}.$$

Let e_k be the step vector supported on one cell. Quantitative Mertens summation and partial summation give, for every positive cell,

$$\sup_{0\le s\le 1}\left|\sum_{t_q\in I_k}\frac{K(s,t_q)}{q}-\int_{I_k}K(s,t)\,\frac{dt}{t}\right| = o_n(1).$$

For the moving low cell, split first at a fixed $\tau>0$. The same estimate applies on $[\tau,\delta]$, whereas $|K(s,t)|\ll st$ makes both the prime sum and the integral over $[t_0,\tau]$ at most $Cs\tau$. Letting first $n\to\infty$ and then $\tau\downarrow 0$ proves the same conclusion for I_0, uniformly in s. Since the fixed positive-cell mesh has finite dimension, summing over the vectors e_k gives input operator-norm convergence for $\|b\|_\infty\le 1$.

The output averaging is a second application of the same argument. Define the actual doubly averaged kernel band operator by

$$(\widehat{B}_n^K b)_i := \frac{1}{H_i}\sum_{p\in\mathcal{P}_i}\frac{1}{p}\sum_{W<q\le y}\frac{K(t_p,t_q)}{q}b_{j(q)}.$$

Equicontinuity in the first variable and Mertens summation now give

$$\sup_{\|b\|_\infty\le 1}\left\|\widehat{B}_n^K b - E_{\mathcal{P}}TJ_{\mathcal{P}}b\right\|_\infty = o_n(1), \tag{8.42}$$

for each fixed permitted mesh. On the low output row the unnormalized kernel numerator is $O(\delta)$, while $H_0\asymp H_0^{(c)}\asymp \log L$; hence the fixed-cutoff Mertens constant is absorbed after normalization. The diagonal multiplier is handled at the same time:

$$\frac{1}{H_i}\sum_{p\in\mathcal{P}_i}\frac{h(t_p)}{p}-\frac{1}{H_i^{(c)}}\int_{I_i}h(t)\,\frac{dt}{t} = o_n(1).$$

This includes I_0, since $h(t)=1+O(t)$ there. Thus both prime indices and the diagonal, rather than only the input prime index, converge to the compressed continuum operator. The operator constants are mesh-uniform, but the convergence threshold may depend on the one fixed mesh. Equation (8.37) handles the remaining error E.

Fix one permitted positive-cell mesh; only its low endpoint now moves with n. Define the *leading* arithmetic band operator, before any marked, Bernoulli, prime-power, guard, or nuisance error is added, by

$$(\widehat{B}_n^{(0)}b)_i := \frac{1}{H_i}\sum_{p\in\mathcal{P}_i}\frac{1}{p}\left\{h(t_p)b_i+\sum_{\substack{W<q\le y\\ q\ne p}}\frac{K(t_p,t_q)}{q}\,b_{j(q)}\right\}.$$

The omitted $p=q$ kernel diagonal has operator norm $o_n(1)$:

$$\max_i \frac{1}{H_i}\sum_{p\in\mathcal{P}_i}\frac{|K(t_p,t_p)|}{p^2} = o_n(1).$$

Indeed $K(t,t)=O(t^2)$; on I_0 the numerator is $O_W((\log y)^{-2})$, and on a positive cell every prime tends to infinity. Consequently the preceding double quadrature, the multiplier comparison, and (8.40) imply

$$\left\|P_{\alpha,n}\widehat{B}_n^{(0)}P_{\alpha,n}-P_{\alpha,c}B_{t_0,\mathcal{P}}P_{\alpha,c}\right\|_{\infty\to\infty} = o_{n;W,\mathcal{P}}(1). \tag{8.43}$$

All operators in this display act on the common coordinate space; the two ranges are still different and are identified next.

Put

$$\Phi_n := P_{\alpha,n}|_{\mathcal{G}_c} : \mathcal{G}_c \longrightarrow \mathcal{G}_n, \qquad \varepsilon_n := \|P_{\alpha,n} - P_{\alpha,c}\|_{\infty\to\infty}.$$

For $x \in \mathcal{G}_c$, $\|\Phi_n x - x\|_\infty \le \varepsilon_n \|x\|_\infty$. Hence, if $\varepsilon_n < 1$, Φ_n is injective; both gauge spaces have codimension one, so it is onto, and

$$\|\Phi_n\| \le 1 + \varepsilon_n, \qquad \|\Phi_n^{-1}\| \le (1-\varepsilon_n)^{-1}. \tag{8.44}$$

If $y = \Phi_n x \in \mathcal{G}_n$, the same calculation gives the ambient-space estimate

$$\|\Phi_n^{-1} y - y\|_\infty = \|x - P_{\alpha,n} x\|_\infty \le \frac{\varepsilon_n}{1-\varepsilon_n}\|y\|_\infty.$$

Pulling the left side of (8.43) back by Φ_n therefore gives the following comparison. Indeed, if $\widehat{T}_n = P_{\alpha,n}\widehat{B}_n^{(0)} P_{\alpha,n}$ and $T_c = P_{\alpha,c} B_{t_0,\mathcal{P}} P_{\alpha,c}$, then for $x \in \mathcal{G}_c$

$$\widehat{T}_n \Phi_n x - \Phi_n T_c x = (\widehat{T}_n - T_c)\Phi_n x + T_c(\Phi_n - I)x + (I - P_{\alpha,n})T_c x.$$

The last two terms are $O(\varepsilon_n)\|x\|_\infty$, because $T_c x \in \mathcal{G}_c$ and the leading operators are uniformly bounded. Together with (8.44), this gives

$$\left\|\Phi_n^{-1}\left(P_{\alpha,n}\widehat{B}_n^{(0)} P_{\alpha,n}\,\big|_{\mathcal{G}_n}\right)\Phi_n - P_{\alpha,c} B_{t_0,\mathcal{P}} P_{\alpha,c}\,\big|_{\mathcal{G}_c}\right\|_{\infty\to\infty} = o_n(1). \tag{8.45}$$

The bounds in (8.34) and a Neumann series now establish the leading arithmetic inverse, with norm at most $2C_c$, for all sufficiently large n. Only *after* this inverse has been obtained do we use (8.36), (8.37), and (8.47): choose W_0 so that their box-independent $O(1/W)$ operator norm, after multiplication by the uniformly bounded projections and gauge maps, is at most $(8C_c)^{-1}$, and then, for the fixed box and mesh, increase n_0 until their vanishing remainders have the same bound. A second Neumann series gives the projected inverse for the actual un-Schur covariance. This order removes any possibility of using an inverse to justify its own leading approximation.

Step 3: quadratic-form convergence on the moving unit sphere. Let $b^T D b = 1$. Then

$$|b_0| \le H_0^{-1/2}, \qquad \sum_{j\ge1} H_j |b_j|^2 \le 1.$$

The norm comparison is now literal. On positive cells the dimension is fixed, $H_j \gg_{\delta,\eta,c_{\mathrm{mesh}}} 1$, and $H_j - H_j^{(c)} = o_n(1)$. On the low cell,

$$|H_0 - H_0^{(c)}|\,|b_0|^2 \ll_W H_0^{-1} = o_n(1).$$

Thus

$$\|J_{\mathcal{P}} b\|^2_{L^2([t_0,1],dt/t)} = 1 + o_n(1).$$

The extension on $(0,t_0)$ has squared norm $|b_0|^2/2 = o_n(1)$, so the full reconstructed norm is $1 + o_n(1)$. The same decomposition, with one factor t, shows that an arithmetic gauge vector has continuum t-pairing $o_n(1)$.

We next compare the leading forms, splitting both prime indices into I_0 and the positive cells. The positive–positive block converges uniformly on the present unit sphere by the fixed finite-dimensional two-index quadrature. For a block with one low index, $|K(s,t)| \le Cst$, Cauchy–Schwarz on the positive cells, and $\sum_{p\in\mathcal{P}_0} t_p/p = O(\delta)$ give

$$|\mathscr{K}_{0,+}(b)| \ll |b_0|\left(\sum_{p\in\mathcal{P}_0}\frac{t_p}{p}\right)\sum_{j\ge1}|b_j|\sum_{q\in\mathcal{P}_j}\frac{t_q}{q} \ll H_0^{-1/2} = o_n(1).$$

Here the last sum is uniformly bounded by Cauchy–Schwarz:

$$\sum_{j\geq 1} H_j\alpha_j|b_j| \leq \left(\sum_{j\geq 1} H_j|b_j|^2\right)^{1/2}\left(\sum_{j\geq 1} H_j\alpha_j^2\right)^{1/2} \ll 1.$$

The identical bound holds for the continuum low–positive block. The two low–low blocks are each

$$O\left(|b_0|^2\left(\sum_{p\in\mathcal{P}_0}\frac{t_p}{p}\right)^2\right)+O\left(|b_0|^2\left(\int_{I_0}dt\right)^2\right)=O(H_0^{-1})=o_n(1).$$

For the multiplier part, write $h=1+(h-1)$. The constant term leaves only $(H_0-H_0^{(c)})b_0^2=o(1)$ on the low row, and the remainder is handled by weighted Mertens summation because $(h(t)-1)/t$ extends continuously to $t=0$. The positive multiplier rows converge by ordinary cellwise quadrature. Finally, the removed discrete $p=q$ kernel diagonal is $o_n(1)$ on the unit sphere by the estimate preceding (8.43). Combining these four estimates with (8.31)– (8.33) proves, uniformly for $b^TDb=1$,

$$b^TD\widehat{B}_n^{(0)}b-\langle\widetilde{f},A\widetilde{f}\rangle_{\mathcal{H}}=o_n(1), \tag{8.46}$$

where $\widetilde{f}$ is the linear low extension of $J_{\mathcal{P}}b$. This explicitly proves both the norm and quadratic-form convergence needed for the Poincaré gap; no equivalence constant depending on H_0 has been used. Indeed, if in addition $b\in\mathcal{G}_n$, then the norm and gauge comparisons above give

$$\|\widetilde{f}\|_{\mathcal{H}}^2=1+o_n(1), \qquad \langle\widetilde{f},t\rangle_{\mathcal{H}}=o_n(1).$$

Since $\|t\|_{\mathcal{H}}^2=1/2$, the exact Hilbert-space projection formula yields

$$\inf_{\lambda\in\mathbb{R}}\|\widetilde{f}-\lambda t\|_{\mathcal{H}}^2=\|\widetilde{f}\|_{\mathcal{H}}^2-\frac{|\langle\widetilde{f},t\rangle_{\mathcal{H}}|^2}{\|t\|_{\mathcal{H}}^2}=1+o_n(1).$$

Thus (8.12) and (8.46) give, for all sufficiently large n,

$$b^TD\widehat{B}_n^{(0)}b\geq\frac{\kappa_0}{2}\,b^TDb \qquad (b\in\mathcal{G}_n).$$

To transfer every remaining arithmetic error literally, let $R=(R_{ij})$ be any symmetric band covariance error satisfying

$$\sup_i\frac{1}{H_i}\sum_j|R_{ij}|\leq\varepsilon.$$

The normalized operator $D^{-1}R$ is self-adjoint for the D-inner product, and

$$\begin{aligned}|b^TRb| &\leq \frac{1}{2}\sum_{i,j}|R_{ij}|(|b_i|^2+|b_j|^2)\\ &=\sum_i|b_i|^2\sum_j|R_{ij}|\\ &\leq\varepsilon\,b^TDb.\end{aligned}$$

Thus its D-Hilbert operator norm is at most ε. Applying this estimate to the squarefree, prime-power, and quadrature errors transfers the continuum quadratic gap to the prime-band covariance form.

Replacing $1_{p|m}$ by $v_p(m)$ adds

$$\sup_p\sum_q p|E_{pq}^{\mathrm{pow}}|\leq\frac{C_{\mathrm{pow}}}{W}+\epsilon_{B,W}(n) \tag{8.47}$$

by (7.56). The constants $C_{\mathrm{sf}}, C_{\mathrm{pow}}$ are known before the compact ODE box is chosen. This is the estimate invoked only after the leading inverse in (8.45) was established. The already chosen $W \geq W_0$ makes its $1/W$ contribution smaller than the reserved continuum inverse and form gaps; after a fixed box is specified, only the threshold for n is increased to absorb the vanishing remainder.

Step 4: the finite nuisance block and its Schur perturbation. It remains to justify the joint Schur complement in (8.23). Every head pattern retains a fixed mass in each of two separated positive-length physical subpools. If a centered linear combination $aR + \sum_\nu b_\nu H_\nu^\circ$ vanished, varying R inside either physical subpool forces $a = 0$, and the affinely independent simplex head patterns then force every $b_\nu = 0$. The fixed mass and separation margins make this quantitative:

$$\Gamma_\xi \geq \gamma_W I \tag{8.48}$$

throughout every subsequently fixed compact tilt box, with γ_W chosen before the radius of that box. We prove this uniformity directly; no convergence of the baseline mixture weights is required. By Lemma 8.1, there is a number $\lambda_W > 0$ such that $\lambda_{e,\sigma}(n) \geq \lambda_W$ for every head pattern e, both signs σ, and all sufficiently large n. Moreover

$$\frac{\inf I_+(n)}{n} - \frac{\sup I_-(n)}{n} = a_+ - b_- > 0,$$

so the corresponding ranges of $R = \log(m/n)$ have a fixed positive separation.

We first prove a finite-n, uniform baseline bound. For $F = aR + \sum_\nu b_\nu H_\nu^\circ$, compare, within each fixed head pattern, its conditional means on the two physical cells. The head term is identical on those cells, whereas their R-means are separated by at least $\log a_+ - \log b_- > 0$. The elementary two-point variance bound and the lower bound λ_W therefore give

$$\mathrm{Var}_{\nu_n}(F) \geq c_{R,W} a^2.$$

On the other hand the finitely many head-pattern vectors are affinely independent after the constant indicator has been deleted. Their masses are all at least $2\lambda_W$, so, uniformly in n,

$$\mathrm{Var}_{\nu_n}\left(\sum_\nu b_\nu H_\nu^\circ\right) \geq c_{H,W} \sum_\nu b_\nu^2.$$

Since R is uniformly bounded,

$$\mathrm{Var}_{\nu_n}\left(\sum_\nu b_\nu H_\nu^\circ\right) \leq 2\,\mathrm{Var}_{\nu_n}(F) + 2a^2\,\mathrm{Var}_{\nu_n}(R),$$

and the preceding a^2-bound absorbs the last term, giving

$$\mathrm{Var}_{\nu_n}\left(\sum_\nu b_\nu H_\nu^\circ\right) \leq C_W\,\mathrm{Var}_{\nu_n}(F).$$

The head-pattern lower bound now yields $\mathrm{Var}_{\nu_n}(F) \geq (c_{H,W}/C_W)\sum_\nu b_\nu^2$. Combining this with the a^2-bound, one may choose $\gamma_W > 0$, independently of n and of every later tilt box, such that

$$\Gamma_{0,n} \geq 2\gamma_W I \qquad (n \geq n_0(W)).$$

It remains only to compare each tilted covariance with its own finite-n baseline. If $S_\xi = \sum_j \eta_j \Omega_j$ is the medium-prime part of the score, then $|S_\xi| \leq B\Omega$, $e^{|S_\xi|/L} \ll_{B,W} 1$, and the one-mark estimates give $\mathbb{E}_{\nu_n}\Omega \ll_W \log L$. Consequently

$$\sup_{\xi \text{ in the fixed } B\text{-box}} \mathbb{E}_{\nu_n}\left|e^{S_\xi/L} - 1\right| \ll_{B,W} \frac{\log L}{L} = o(1).$$

The physical and finite-head factors are $1 + O_{B,W}(L^{-1})$, and normalization changes the same estimate by only a constant factor. Thus

$$\sup_{\xi \text{ in the fixed } B\text{-box}} \|\nu_{n,\xi} - \nu_n\|_{\mathrm{TV}} = o_{n;B,W}(1).$$

Every coordinate of Z is uniformly bounded and its dimension depends only on W. If $d_W = \dim Z$ and $\max_\nu \|Z_\nu\|_\infty \le M_W$, then for every Euclidean unit vector a,

$$\left|\mathrm{Var}_{\nu_{n,\xi}}(a^T Z) - \mathrm{Var}_{\nu_n}(a^T Z)\right| \le 6 d_W M_W^2 \|\nu_{n,\xi} - \nu_n\|_{\mathrm{TV}}.$$

Taking the supremum over a gives

$$\sup_{\xi \text{ in the fixed } B\text{-box}} \|\Gamma_{\xi,n} - \Gamma_{0,n}\|_{\mathrm{op}} = o_{n;B,W}(1).$$

For each fixed B, increasing only $n_0(B,W)$ now proves (8.48). Its constant γ_W was chosen from the uniform finite-n geometry before B, so no unproved limiting mixture and no box-dependent main constant is involved. The constant score is absent: normalization has removed it, and one centered head indicator was deleted before Γ_ξ was formed. The head cross rows are $O_{B,W}(1/L)$ throughout the fixed box, and the physical cross row is given by the Stieltjes estimate (7.61). Write

$$U_\xi := \mathrm{Cov}_\xi(\Omega, Z), \qquad C_\xi := \mathrm{Cov}_\xi(\Omega, \Omega).$$

The actual block elimination is

$$q_n^{-1} \mathcal{C}_\xi^Z = C_\xi - U_\xi \Gamma_\xi^{-1} U_\xi^T.$$

The marked head/physical bounds give, for $\|b\|_\infty \le 1$,

$$|(U_\xi^T b)_\nu| \le \sum_j |(U_\xi)_{j\nu}|\,|b_j| \ll_{B,W} \frac{1}{L} \sum_j H_j \ll_{B,W} \frac{\log L}{L}.$$

Together with $\|\Gamma_\xi^{-1}\|_{\mathrm{op}} \le \gamma_W^{-1}$ this gives, row by row,

$$\left\|D^{-1} U_\xi \Gamma_\xi^{-1} U_\xi^T b\right\|_\infty \ll_{B,W} \frac{\log L}{L^2}, \tag{8.49}$$

because $|(U_\xi)_{i\nu}| \ll_{B,W} H_i/L$ and $\sum_j H_j \ll \log L$. This is the promised normalized row estimate, with an explicit rate.

Step 5: the exact physical direction and the full quotient form. The physical column is nevertheless responsible for the exact null relation, and we record it rather than discarding it as a small error. Let

$$T_y := \sum_{W < p \le y} t_p v_p, \qquad S_\alpha := \sum_j \alpha_j \Omega_j, \qquad S_g := \sum_{W < p \le y} (\alpha_{j(p)} - t_p) v_p.$$

For a statistic X, put $X^{\circ,\xi} := X - \mathbb{E}_{\nu_{n,\xi}} X$. This notation is distinct from the already defined baseline-centered head scores H_e°; changing their centering subtracts only a constant and therefore leaves the centered covariance span unchanged. Every member of every active structured cell is y-smooth, so its logarithm is the sum of the $p \le y$ valuation terms. Moreover, for each $p \le W$, the function v_p is constant on every exact head-pattern cell and hence is a linear combination of the retained head indicators and the deleted constant indicator. After centering, the constant disappears and the whole $p \le W$ term lies in the centered head span. Therefore, identically at finite n,

$$T_y^{\circ,\xi} = \frac{R^{\circ,\xi} - \sum_{p \le W} (\log p) v_p^{\circ,\xi}}{\log y} \in \operatorname{span} Z^{\circ,\xi}, \qquad S_\alpha^{\circ,\xi} = T_y^{\circ,\xi} + S_g^{\circ,\xi}. \tag{8.50}$$

Let $P_{Z^\perp}$ denote orthogonal projection in $L^2(\nu_{n,\xi})$ onto the orthogonal complement of the centered coordinates of Z. Thus $P_{Z^\perp}$ kills $T_y^{\circ,\xi}$ exactly. More generally, for $b = q + \lambda\alpha$, with $q \in \mathcal{G}_n$, and for every $\mu \in \mathbb{R}$,

$$P_{Z^\perp}\{(b \cdot \Omega)^{\circ,\xi}\} = P_{Z^\perp}\left\{\left(\sum_{W<p\le y}(b_{j(p)} - \mu t_p)v_p\right)^{\circ,\xi}\right\}. \tag{8.51}$$

This is the precise finite-n post-Schur relation used below. We next record its arithmetic geometry. Define

$$V_n := \sum_{W<p\le y}\frac{(\alpha_{j(p)} - t_p)^2}{p}, \qquad \mathfrak{d}_n(b)^2 := \inf_{\mu\in\mathbb{R}}\sum_{W<p\le y}\frac{|b_{j(p)} - \mu t_p|^2}{p}, \qquad A_{\alpha,n} := \alpha^T D\alpha.$$

Since $\sum_{p\in\mathcal{P}_j}(\alpha_j - t_p)/p = 0$ in every band, and since $q^T D\alpha = 0$, the following identities are exact. For all sufficiently large n, every permitted band has positive harmonic mass and $\alpha_j > 0$; hence $A_{\alpha,n} > 0$, while $V_n \ge 0$. In particular the denominator in the third line below is nonzero:

$$\begin{aligned}
\mathfrak{d}_n(q+\lambda\alpha)^2 &= \inf_{\mu\in\mathbb{R}}\left\{\|q + (\lambda-\mu)\alpha\|_D^2 + \mu^2 V_n\right\} \\
&= \|q\|_D^2 + \inf_{\mu\in\mathbb{R}}\left\{(\lambda-\mu)^2 A_{\alpha,n} + \mu^2 V_n\right\} \\
&= \|q\|_D^2 + \lambda^2\frac{A_{\alpha,n}V_n}{A_{\alpha,n} + V_n} \ge \|q\|_D^2.
\end{aligned} \tag{8.52}$$

We also record here the size of V_n, independently of the later regression argument. Weighted Mertens summation and (8.40) give

$$V_n = \sum_j \int_{I_j} |t - \alpha_j^{(c)}|^2\,\frac{dt}{t} + o_n((\delta+\eta)^2).$$

On $I_0 = [t_0, \delta]$, the displayed integral is exactly

$$\frac{\delta^2 - t_0^2}{2} - \frac{(\delta - t_0)^2}{\log(\delta/t_0)} \asymp \delta^2.$$

On a positive cell $I = [a, a+\ell]$, it is $\asymp \ell^3/a$; summing the regular relative mesh gives $\asymp \eta^2$. Consequently, uniformly over permitted meshes,

$$V_n \asymp \delta^2 + \eta^2 \asymp (\delta+\eta)^2.$$

We now transfer this exact distance to the covariance form; this step is needed because an arithmetic center is not a continuum center. If $\mathfrak{d}_n(b) > 0$, let μ_* be its unique minimizer and put

$$c_p := b_{j(p)} - \mu_* t_p. \tag{8.53}$$

Differentiation of the finite sum gives the exact normal equation

$$\sum_{W<p\le y}\frac{t_p c_p}{p} = 0. \tag{8.54}$$

Normalize for the moment so that $\mathfrak{d}_n(b) = 1$. Equation (8.52), together with $A_{\alpha,n} \asymp 1$ and $V_n \asymp (\delta+\eta)^2$, gives

$$\|q\|_D \le 1, \qquad |\lambda| + |\mu_*| \ll_{\delta,\eta} 1.$$

In particular $|q_0| \le H_0^{-1/2}$, $\alpha_0 \ll_\delta H_0^{-1}$, and $t_0 \to 0$.

Define the piecewise affine continuum coefficient

$$f_n(t) := b_j - \mu_* t \qquad (t \in I_j),$$

and extend it on $(0, t_0)$ by $f_n(t_0)t/t_0$; denote the resulting function by $\widetilde{f}_n$. The cellwise Mertens argument in (8.38)– (8.39) gives

$$\|\widetilde{f}_n\|_{\mathcal{H}}^2 = 1 + o_n(1), \qquad \langle \widetilde{f}_n, t\rangle_{\mathcal{H}} = o_n(1). \tag{8.55}$$

Here the second relation is (8.54) after quadrature. For clarity about the moving low cell, the only term that can retain the fixed-cutoff Mertens constant is $|b_0|^2(H_0 - H_0^{(c)}) = O_W(|b_0|^2) = o_n(1)$; every term carrying a factor t converges by weighted Mertens summation. Also $|f_n(t_0)| = o_n(1)$, so the linear extension has $o_n(1)$ squared norm and $o_n(1)$ t-pairing. Thus no convergence $D \to D_c$ has been assumed.

For prime coefficients $c = (c_p)$, write the leading discrete form as

$$\mathscr{Q}_n^{(0)}(c) := \sum_{W<p\le y} \frac{h(t_p)c_p^2}{p} + \sum_{\substack{W<p,r\le y\\ p\ne r}} \frac{K(t_p,t_r)c_pc_r}{pr}.$$

The continuum double integral contains its diagonal as a measure-zero set, whereas the finite prime double sum in the quadrature temporarily contains $p = r$. The difference is

$$\Delta_{\mathrm{diag}}(c) := \sum_{W<p\le y} \frac{K(t_p,t_p)c_p^2}{p^2} = o_n(1) \qquad (\mathfrak{d}_n(b) = 1).$$

To verify this uniformly, on the positive cells $p \ge y^\delta$, so

$$\sum_{t_p>\delta} \frac{|K(t_p,t_p)|c_p^2}{p^2} \ll y^{-\delta} \sum_{t_p>\delta} \frac{c_p^2}{p} = o(1).$$

On I_0, $c_p = b_0 - \mu_* t_p$, where $|b_0| = O(H_0^{-1/2})$ and $\mu_* = O_{\delta,\eta}(1)$. Since

$$\sum_{p>W} \frac{t_p^2}{p^2} = O_W((\log y)^{-2}), \qquad \sum_{p>W} \frac{t_p^4}{p^2} = O_W((\log y)^{-4}),$$

and $K(t,t) = O(t^2)$, the low-cell contribution is $o(1)$ as well. The two-index quadrature proved in (8.42)– (8.46) applies also to $f_n = J_{\mathcal{P}}b - \mu_* t$: the step part is covered by the displayed operator comparison, while the t-part follows from the same Stieltjes summation, and its coefficient is bounded under the present normalization. Together with (8.55), it gives

$$\mathscr{Q}_n^{(0)}(c) = \langle \widetilde{f}_n, A\widetilde{f}_n\rangle_{\mathcal{H}} + o_n(1).$$

All $o_n(1)$ terms in these comparisons are uniform on the normalized class $\mathfrak{d}_n(b) = 1$: q is bounded in D-norm, $\lambda, \mu_* = O_{\delta,\eta}(1)$, the positive-cell dimension is fixed, and the low coordinate tends to zero. In particular, the resulting threshold for n does not depend on b. The Poincaré inequality (8.12) now applies: the two relations in (8.55) say that the squared distance of $\widetilde{f}_n$ from span$\{t\}$ is $1 + o_n(1)$. Hence, by homogeneity, there is a mesh-uniform $\kappa_1 > 0$ such that, for every b and all sufficiently large n,

$$\mathscr{Q}_n^{(0)}(c) \ge \kappa_1 \mathfrak{d}_n(b)^2 \tag{8.56}$$

when c is defined by (8.53). If $\mathfrak{d}_n(b) = 0$, then (8.52) gives $b = 0$, so the assertion is trivial.

It remains to verify that every arithmetic error is relative to this same minimizing distance. If $E_{pr} = E_{rp}$ is any prime covariance error with $\sup_p \sum_r p|E_{pr}| \le \varepsilon$, symmetry and $2|c_p c_r| \le c_p^2 + c_r^2$ give the exact estimate

$$\left|\sum_{p,r} E_{pr} c_p c_r\right| \le \varepsilon \sum_p \frac{c_p^2}{p} = \varepsilon \mathfrak{d}_n(b)^2. \tag{8.57}$$

The proof of Corollary 7.6, before band averaging, gives the corresponding prime-level guard estimate

$$\sup_{W<p\le y} \sum_{W<r\le y} p|E_{pr}^{\mathrm{guard}}| \ll_{B,W} \frac{y^2 L^{A+3}}{n} = o_n(1)$$

for the fixed exponent A in that corollary. Explicitly, a deleted integer satisfies the two score envelopes (7.68). Taking the signs of the covariance differences as the coefficient vector b in (7.69), the exact conditional covariance identity (7.67) gives $O_{B,W}(y^2 L^{A+2}/n)$, which is stronger than the displayed bound. Hence (8.57) applies to guarded deletions as well as to the analytic covariance remainders. Therefore the squarefree and prime-power remainders in (8.36) and (8.47), including the guarded-deletion remainder from Corollary 7.6, contribute at most

$$\left\{\frac{C_{\mathrm{sf}} + C_{\mathrm{pow}}}{W} + o_{n;B,W,\mathcal{P}}(1)\right\} \mathfrak{d}_n(b)^2.$$

The main constants here are independent of the subsequently fixed tilt box. Equivalently, if $\mathscr{E}_n(c)$ denotes the sum of these actual covariance errors, then

$$\operatorname{Var}_\xi\left(\sum_p c_p v_p\right) = \mathscr{Q}_n^{(0)}(c) + \mathscr{E}_n(c), \qquad |\mathscr{E}_n(c)| \le \left\{\frac{C_{\mathrm{sf}} + C_{\mathrm{pow}}}{W} + o_{n;B,W,\mathcal{P}}(1)\right\} \mathfrak{d}_n(b)^2.$$

Finally, the one-mark head estimates and (7.61) give

$$\left\|\operatorname{Cov}_\xi\left(Z, \sum_p c_p v_p\right)\right\| \ll_{B,W} \frac{1}{L} \sum_p \frac{|c_p|}{p} \ll_{B,W} \frac{\sqrt{\log L}}{L} \mathfrak{d}_n(b).$$

Together with (8.48), the nuisance Schur subtraction is consequently

$$O_{B,W}\left(\frac{\log L}{L^2}\right) \mathfrak{d}_n(b)^2.$$

Indeed, the definition (8.23) and the exact score identity (8.51) give

$$\frac{1}{q_n} b^T \mathcal{C}_\xi^Z b = \left\|P_{Z^\perp}\{(b\cdot\Omega)^{\circ,\xi}\}\right\|_2^2 = \left\|P_{Z^\perp}\left\{\left(\sum_p c_p v_p\right)^{\circ,\xi}\right\}\right\|_2^2.$$

First choose W so that the box-independent $1/W$ terms are smaller than half the Poincaré gap, and then take n sufficiently large for the fixed box and mesh. Equations (8.56) and (8.52) then yield

$$\frac{1}{q_n} b^T \mathcal{C}_\xi^Z b \ge \kappa\, \mathfrak{d}_n(b)^2 \ge \kappa \|q\|_D^2.$$

Every band vector has a unique decomposition $b = q + \lambda\alpha$, $q \in \mathcal{G}_n$, and $\|q\|_D^2 = \inf_\lambda \|b - \lambda\alpha\|_D^2$. Thus the last display is exactly (8.25); no gap on $\mathcal{G}_n$ has been used as a substitute for the full quotient inequality.

The inverse assertion is a separate projected statement. The gauge isomorphism (8.44), the pulled comparison (8.45), and the two successive Neumann steps following that display give the projected ℓ^∞-inverse for the actual un-Schur band covariance C_ξ. Its norm is bounded independently of the later box and mesh; only the threshold for the vanishing arithmetic remainder depends on them. Apply Lemma 8.3 with $P = P_{\alpha,n}$ and $E = U_\xi \Gamma_\xi^{-1} U_\xi^T$. Equation (8.49) and the symmetric band-row calculation preceding (8.47) give the required operator and quadratic perturbation bounds $O_{B,W}(\log L/L^2) = o_{n;B,W}(1)$. The lemma therefore preserves the projected inverse, proving (8.26). This proves (8.25)– (8.26) with the constant row removed and the head/physical block jointly, rather than separately, inverted. □

8.5 The weighted compensated score

The quotient estimate alone is not sufficient: the scalar parameter which remains after quotient fitting is of order $1/(\delta+\eta)$. We now prove the weighted estimate which shows that its *effective* prime coefficients remain bounded.

Put

$$w := \delta + \eta, \qquad g_p := \alpha_{j(p)} - t_p, \qquad V := V_n = \sum_{W<p\le y} \frac{g_p^2}{p}.$$

The regular lower bound in (8.20) is used in the next lemma.

For a band vector put

$$\|b\|_\sharp := \max_j \frac{|b_j|}{\alpha_j}.$$

This is a finite norm for every n, but its low weight $\alpha_0 \asymp \delta/H_0$ moves with n. For the continuum cells we use the distinct notation

$$\|b\|_{\sharp,c} := \max_j \frac{|b_j|}{\alpha_j^{(c)}}.$$

Lemma 8.5 (Weighted discrete quotient inverse)**.** *Assume that $W \ge W_0$ and $w = \delta + \eta \le w_0$ have been chosen as in Lemma 8.4. Let*

$$\mathsf{B}_\xi^Z := D^{-1}\mathcal{C}_\xi^Z / q_n, \qquad \mathcal{G} := \{b : \alpha^T D b = 0\}.$$

If $u \in \mathcal{G}$, $|u_j| \le Cw\alpha_j$, and the gauge-fixed solution $q \in \mathcal{G}$ is defined by

$$P_{\alpha,n}\mathsf{B}_\xi^Z P_{\alpha,n} q = u, \tag{8.58}$$

then

$$|q_j| \le C_\sharp w \alpha_j \quad (0 \le j \le s). \tag{8.59}$$

Here $C_\sharp$ depends on the displayed right-side constant C and the continuum operator, but not on W, the compact tilt radius, or the permitted fine mesh once $W \ge W_0$. This is uniform in the moving cutoff, every permitted regular mesh, all sufficiently large n, and every fixed compact effective-tilt box, in the precise sense that the constant is independent of the mesh and box radius, while the threshold for n may depend on $(B, W, \mathcal{P})$.

Proof. Step 1: a sharp continuum inverse. We first prove the continuum assertion, using only the continuum gauge $\mathcal{G}_c$, centers $\alpha^{(c)}$, and harmonic matrix D_c. For a band vector define the linear reconstruction

$$(J_\mathcal{P}^\sharp q)(t) := \frac{t}{\alpha_j^{(c)}} q_j \quad (t \in I_j).$$

Thus

$$\|J_\mathcal{P}^\sharp q\|_{X_t} = \|q\|_{\sharp,c},$$

including on the moving low cell. If a uniform continuum inverse did not exist, there would be a sequence of moving cutoffs and permitted meshes, numbers $\varepsilon_\nu \downarrow 0$, and $q_\nu \in \mathcal{G}_{c,\nu}$ such that

$$\|q_\nu\|_{\sharp,c} = 1, \qquad \max_j \frac{|[P_{\alpha,c} B_{t_0,\mathcal{P}} q_\nu]_j|}{\alpha_j^{(c)}} \le \varepsilon_\nu.$$

As in (8.28), there is a bounded scalar β_ν such that

$$(B_{t_0,\mathcal{P}}q_\nu)_j = \beta_\nu \alpha_j^{(c)} + r_{\nu,j}, \qquad |r_{\nu,j}| \le \varepsilon_\nu \alpha_j^{(c)}. \tag{8.60}$$

Here one may take

$$r_\nu = P_{\alpha,c} B_{t_0,\mathcal{P}} q_\nu, \qquad \beta_\nu = \frac{(\alpha^{(c)})^T D_c B_{t_0,\mathcal{P}} q_\nu}{(\alpha^{(c)})^T D_c \alpha^{(c)}}.$$

The denominator is bounded below uniformly by the cells meeting $[1/2,1]$. Also $\|q_\nu\|_\infty \le \|q_\nu\|_{\sharp,c} = 1$, $\sum_j H_j^{(c)}\alpha_j^{(c)} = 1-t_0$, and the compressed continuum operators are uniformly bounded. Hence $|\beta_\nu| \le C$, with C independent of the moving cutoff and mesh. For the local estimates below write $q = q_\nu$, $\beta = \beta_\nu$, and suppress ν from the cell data. Put $x_j = q_j/\alpha_j^{(c)}$ and, on I_j, set

$$d(t) := (J_{\mathcal{P}}^\sharp - J_{\mathcal{P}})q(t) = x_j(t-\alpha_j^{(c)}).$$

Then $\|x\|_\infty = \|q\|_{\sharp,c}$ and, exactly on every cell,

$$\int_{I_j} d(t)\,\frac{dt}{t} = 0.$$

For the multiplier part the constant value $h(\alpha_j^{(c)})$ therefore cancels. On a positive relative cell, Lipschitz continuity of h, the mesh bounds, and $\alpha_j^{(c)} \asymp \inf I_j$ give

$$\frac{1}{H_j^{(c)}}\left|\int_{I_j} h(t)d(t)\,\frac{dt}{t}\right| \le C\eta\alpha_j^{(c)}\|q\|_{\sharp,c}.$$

On I_0, use $h(t) = 1 + O(t)$ and the two explicit estimates

$$\frac{1}{H_0^{(c)}}\int_{I_0}(t-\alpha_0^{(c)})\,\frac{dt}{t} = 0, \qquad \frac{1}{H_0^{(c)}}\int_{I_0}|t-\alpha_0^{(c)}|\,\frac{dt}{t} \le 2\alpha_0^{(c)}.$$

They give an $O(\delta\alpha_0^{(c)})\|q\|_{\sharp,c}$ low multiplier error. For the integral part, $|K(s,t)| \le Cst$ gives

$$\left|\frac{1}{H_i^{(c)}}\int_{I_i}\int_{t_0}^1 K(s,t)d(t)\,\frac{dt}{t}\,\frac{ds}{s}\right| \le C\alpha_i^{(c)}\|q\|_{\sharp,c}\sum_k \int_{I_k}|t-\alpha_k^{(c)}|\,dt.$$

The low summand is $O(\delta^2)$. On the regular positive cells the sum is $O(\eta)$: each summand is $O(\eta^2(\inf I_k)^2)$, and comparison with the corresponding geometric Riemann sum gives $O(\eta)$. Hence, uniformly over all permitted meshes,

$$\max_j \frac{\left|[E_{\mathcal{P}}A^{(t_0)}(J_{\mathcal{P}}^\sharp - J_{\mathcal{P}})q]_j\right|}{\alpha_j^{(c)}} \le Cw\|q\|_{\sharp,c}.$$

If $G(s) = A^{(t_0)}J_{\mathcal{P}}^\sharp q(s)/s$, then (8.19), $h \in C^1$, and the fact that $J_{\mathcal{P}}^\sharp q(t)/t = x_j$ is constant on each cell give

$$\operatorname{osc}_{I_0} G \le C\delta\|q\|_{\sharp,c}, \qquad \operatorname{osc}_{I_j} G \le C\eta\|q\|_{\sharp,c} \quad (j \ge 1).$$

Moreover, since $H_j^{(c)}\alpha_j^{(c)} = |I_j|$, for every cell

$$(E_{\mathcal{P}}A^{(t_0)}J_{\mathcal{P}}^\sharp q)_j = \alpha_j^{(c)}\frac{1}{|I_j|}\int_{I_j} G(s)\,ds.$$

In particular, on the low cell,

$$(E_{\mathcal{P}} A^{(t_0)} J^{\sharp}_{\mathcal{P}} q)_0 = \alpha_0^{(c)} \frac{1}{\delta - t_0} \int_{t_0}^{\delta} G(s)\, ds.$$

Consequently (8.60) implies

$$\left\| \frac{A^{(t_0)} J^{\sharp}_{\mathcal{P}} q}{t} - \beta \right\|_{\infty} \ll w + o(1). \tag{8.61}$$

Extend the low linear formula to $(0, t_0)$, and denote the resulting function by $\widetilde{f}$. Since $|K(s,t)/s| \ll t$, the omitted input changes the weighted output by $O(t_0^2)$. For the extended function, its quotient by t is constant on $(0, \delta]$. The formulas for h and $K(s,t)/s$, with $F \in C^2[0,2]$, therefore give

$$\operatorname{osc}_{0 \le s \le \delta} \frac{A\widetilde{f}(s)}{s} = O(\delta).$$

At $s = 0$, the quotient kernel is interpreted by (8.19). Hence (8.61) holds on all of $(0,1]$ for the full operator, with the same $O(w) + o(1)$ right side. Put $e(t) = A\widetilde{f}(t)/t - \beta$. Self-adjointness and $At = 0$ give the exact identity

$$0 = \langle t, A\widetilde{f} \rangle_{\mathcal{H}} = \frac{\beta}{2} + \int_0^1 t\, e(t)\, dt.$$

Hence $|\beta| \le \|e\|_\infty = O(w) + o(1)$. The X_t inverse (8.14) then puts the reconstruction within $O(w) + o(1)$ of λt. Since $(\alpha^{(c)})^T D_c q = 0$, this means exactly

$$0 = \sum_j H_j^{(c)} (\alpha_j^{(c)})^2 \frac{q_j}{\alpha_j^{(c)}}.$$

The total weight in this sum is bounded below by a structural positive constant: the cells contained in $[1/2, 1]$ have centers at least $1/2$ and total harmonic mass bounded below. Averaging $q_j/\alpha_j^{(c)} = \lambda + O(w) + o(1)$ against these positive weights therefore gives $\lambda = O(w) + o(1)$. Once the structural threshold w_0 is sufficiently small, this contradicts $\|q\|_{\sharp,c} = 1$. Equivalently, with

$$\Pi_c := S_c^{-1} P_{\alpha,c} S_c, \qquad \mathcal{R}_c := \operatorname{Ran} \Pi_c, \qquad \mathcal{T}_c := \Pi_c S_c^{-1} B_{t_0,\mathcal{P}} S_c \Pi_c,$$

there is a mesh- and t_0-uniform constant $C_\sharp^{(0)}$ such that

$$\left\| (\mathcal{T}_c|_{\mathcal{R}_c})^{-1} \right\|_{\infty \to \infty} \le C_\sharp^{(0)}. \tag{8.62}$$

Step 2: arithmetic cell data and sharp h, K quadrature. The arithmetic transfer must be relative. If $\|b\|_\sharp \le 1$, then $|b_{j(q)}| \le \alpha_{j(q)} \ll t_q$ on positive cells, while on the low cell $|b_0| \le \alpha_0$. The pointwise form of (7.14) and (7.10) is, for $p \ne q$,

$$p|E^{\mathrm{sf}}_{pq}| \ll_{B,W} \frac{1}{qL}. \tag{8.63}$$

The diagonal has the same L^{-1} error, in addition to the Bernoulli correction treated below. Therefore the exact band average (8.37) gives

$$\begin{aligned} |(\mathsf{E}_n E^{\mathrm{sf}} J_n b)_0| &\ll_{B,W} L^{-1} \ll_{B,W} \alpha_0 \frac{\log L}{\delta L}, \\ |(\mathsf{E}_n E^{\mathrm{sf}} J_n b)_i| &\ll_{B,W} \alpha_i \frac{1}{\delta L} \quad (i \ge 1), \\ \max_i \frac{|(\mathsf{E}_n E^{\mathrm{sf}} J_n b)_i|}{\alpha_i} &\ll_{B,W} \frac{\log L}{\delta L}. \end{aligned}$$

For the low row, the first inequality is the explicit calculation

$$\frac{1}{H_0}\sum_{p\in\mathcal{P}_0}\frac{1}{p}\sum_{q}O_{B,W}\left(\frac{|b_{j(q)}|}{qL}\right)\ll_{B,W}\frac{1}{L},$$

because $\sum_q |b_{j(q)}|/q \ll H_0\alpha_0+\sum_{j\geq 1}H_j\alpha_j \ll 1$. Thus the numerator is genuinely $o(\alpha_0)$, since $\alpha_0 \asymp \delta/\log L$. The active cells have the numerical guards deleted. By (7.63), that deletion contributes an operator G_n satisfying, for $\|b\|_\sharp \leq 1$,

$$\max_i \frac{|(G_n b)_i|}{\alpha_i}=o_{n;B,W,\mathcal{P}}(1),$$

because its band output is $o_n(\alpha_0)$, $\|b\|_\infty\leq 1$, and $\alpha_i \gg \alpha_0$ on every positive cell. Thus guard deletion is part of the vanishing sharp remainder.

We next state the sharp quadrature which replaces an unspecified additive discretization error. Let $\widehat{B}_n^{h,K}$ be the band operator with no marked error, defined by

$$(\widehat{B}_n^{h,K}b)_i := \frac{1}{H_i}\sum_{p\in\mathcal{P}_i}\frac{1}{p}\left\{h(t_p)b_i+\sum_{\substack{W<q\leq y\\ q\neq p}}\frac{K(t_p,t_q)}{q}b_{j(q)}\right\}.$$

Then, for every fixed permitted mesh,

$$\left\|S_n^{-1}\widehat{B}_n^{h,K}S_n - S_c^{-1}B_{t_0,\mathcal{P}}S_c\right\|_{\infty\to\infty}=o_{n;W,\mathcal{P}}(1). \tag{8.64}$$

Here is the low-cell calculation needed for this relative statement. Since $F\in C^2[0,2]$ and $F(0)=1$, the function

$$\varkappa(s,t):=\frac{K(s,t)}{st}$$

has a continuous extension to $[0,1]^2$. The (i,k)-entry of the conjugated discrete kernel, with the diagonal temporarily inserted, is

$$\frac{\alpha_k}{H_i\alpha_i}\sum_{p\in\mathcal{P}_i}\frac{t_p}{p}\sum_{q\in\mathcal{P}_k}\frac{t_q}{q}\varkappa(t_p,t_q). \tag{8.65}$$

To see every normalization, introduce the probability measures

$$\rho_{j,n}:=\frac{1}{H_j\alpha_j}\sum_{p\in\mathcal{P}_j}\frac{t_p}{p}\,\delta_{t_p}, \qquad\qquad \rho_{j,c}:=\frac{1}{H_j^{(c)}\alpha_j^{(c)}}\mathbf{1}_{I_j}(t)\,dt,$$

and retain $\omega_{n,j}=H_j\alpha_j^2$, $\omega_{c,j}=H_j^{(c)}(\alpha_j^{(c)})^2$. Because $H_j\alpha_j=\sum_{p\in\mathcal{P}_j}t_p/p$ and $H_j^{(c)}\alpha_j^{(c)}=|I_j|$, these are indeed probability measures. Formula (8.65) is exactly

$$\omega_{n,k}\iint \varkappa(s,t)\,d\rho_{i,n}(s)d\rho_{k,n}(t), \tag{8.66}$$

and its continuum counterpart is the same expression with n replaced by c.

The moving-low convergence follows from the uniform weighted prime count

$$\sup_{t_0\leq x\leq 1}\left|\sum_{W<p\leq y^x}\frac{t_p}{p}-(x-t_0)\right|=O_W((\log y)^{-1}). \tag{8.67}$$

This is $\sum_{p\le X}(\log p)/p = \log X + O(1)$, divided by $\log y$. Stieltjes summation and uniform continuity of $\varkappa$ therefore give $\rho_{j,n} \Rightarrow \rho_{j,c}$, uniformly against all sections of $\varkappa$, on every positive cell and also on $I_0 = [t_0, \delta]$. On the low cell,

$$\omega_{n,0} = (H_0\alpha_0)\alpha_0 = o_n(1), \qquad \omega_{c,0} = (\delta - t_0)\alpha_0^{(c)} = o_n(1),$$

whereas $\omega_{n,k} - \omega_{c,k} = o_n(1)$ on every positive cell. Thus (8.66) converges in all four low/positive orientations: a low input carries the vanishing factor $\omega_{0,\bullet}$, and a low output is averaged against a probability measure. In particular neither H_0 nor α_0^{-1} is lost. Removing the inserted diagonal changes the conjugated operator by at most

$$\max_i \frac{1}{H_i} \sum_{p\in\mathcal{P}_i} \frac{|K(t_p,t_p)|}{p^2} = o_{n;W,\mathcal{P}}(1).$$

Indeed $|K(t,t)| \ll t^2$. On a positive cell its primes tend to infinity, while on the low cell $t_p^2 = (\log p)^2/(\log y)^2$ and $\sum_{p>W}(\log p)^2/p^2 < \infty$. For the diagonal multiplier on I_0, write $h(t) = 1 + t\widetilde{h}(t)$, where $\widetilde{h}$ is continuous. The constant part has average exactly one in both the arithmetic and continuum rows; (8.67) handles the remaining weighted average, together with $H_0/H_0^{(c)} \to 1$. Ordinary Mertens quadrature applies on every positive cell. The dimension is fixed after the mesh is chosen, so the entrywise convergence is exactly (8.64). Together with (8.41), this also verifies the sharp convergence of the arithmetic and continuum gauges; no convergence of D to D_c is used.

Step 3: the remaining sharp arithmetic errors and the nuisance Schur block. The Bernoulli diagonal correction, kept out of E_{pq}^{sf} in the preceding off-diagonal calculation, satisfies

$$|E_i^{\mathrm{sf,diag}}| \ll \frac{1}{H_i} \sum_{p\in\mathcal{P}_i} \frac{|b_{j(p)}|}{p^2} \le \frac{\alpha_i}{W}.$$

For the prime-power part we use the product-weighted aggregate bounds, rather than an unavailable unweighted comparison of $H_0^{-1}\sum_{p\in\mathcal{P}_0} p^{-2}$ with α_0. For $\|b\|_\sharp \le 1$, harmonic centering and the regular mesh give

$$\sum_q \frac{t_q|b_{j(q)}|}{q} \ll 1, \qquad \sum_q \frac{t_q|b_{j(q)}|}{q^2} \ll \frac{1}{W}.$$

Indeed, on a positive cell $\alpha_j \asymp t_q$, while on the low cell

$$|b_0| \sum_{q\in\mathcal{P}_0} \frac{t_q}{q} \le H_0\alpha_0^2 \ll 1.$$

The unweighted vanishing remainders in the three aggregate estimates are also harmless at the moving low cell. More precisely, for the fixed regular mesh,

$$\sum_q \frac{|b_{j(q)}|}{q^2} \ll_W \alpha_0 + o(\alpha_0), \qquad \frac{1}{H_i\alpha_i} \sum_{p\in\mathcal{P}_i} \frac{1}{p^2} = O_{W,\delta,\eta}(1), \tag{8.68}$$

uniformly in i. For the first estimate, the low-cell contribution is $|b_0|\sum_{p>W} p^{-2} \ll_W \alpha_0$, while every positive cell starts at a power of y and is $o(\alpha_0)$; the second estimate is the same reciprocal calculation after division by the cell's harmonic mass. Also $\sum_q |b_{j(q)}|/q \ll 1$, as in the squarefree calculation above. Thus the unweighted error coefficient in (7.52) remains $o(\alpha_i)$ in each orientation. Denote by $E_i^{JI}, E_i^{IJ}, E_i^{JJ}$, and E_i^{diag} the four prime-power contributions to the normalized band output. We write every orientation rather than invoke symmetry. The box-independent main terms satisfy

$$\begin{aligned}
|E_i^{JI}|_{\mathrm{main}} &\ll \frac{1}{H_i}\sum_{p\in\mathcal{P}_i}\frac{t_p}{p^2}\sum_q \frac{t_q|b_{j(q)}|}{q} \le \frac{\alpha_i}{W}\sum_q \frac{t_q|b_{j(q)}|}{q} \ll \frac{\alpha_i}{W},\\
|E_i^{IJ}|_{\mathrm{main}} &\ll \frac{1}{H_i}\sum_{p\in\mathcal{P}_i}\frac{t_p}{p}\sum_q \frac{t_q|b_{j(q)}|}{q^2} \ll \frac{\alpha_i}{W},\\
|E_i^{JJ}|_{\mathrm{main}} &\ll \frac{1}{H_i}\sum_{p\in\mathcal{P}_i}\frac{t_p}{p^2}\sum_q \frac{t_q|b_{j(q)}|}{q^2} \ll \frac{\alpha_i}{W^2},\\
|E_i^{\mathrm{diag}}|_{\mathrm{main}} &\ll \frac{|b_i|}{H_i}\sum_{p\in\mathcal{P}_i}\frac{1}{p^2} \le \frac{\alpha_i}{W}.
\end{aligned}$$

Let $\epsilon^{(0)}_{B,W}(n)$ be the single pair remainder in (7.52). The four unweighted remainders, divided by α_i, are bounded respectively by

$$\begin{aligned}
&\frac{\epsilon^{(0)}_{B,W}(n)}{H_i\alpha_i}\sum_{p\in\mathcal{P}_i}\frac{1}{p^2}\sum_q\frac{|b_{j(q)}|}{q},\\
&\frac{\epsilon^{(0)}_{B,W}(n)}{\alpha_i}\sum_q\frac{|b_{j(q)}|}{q^2},\\
&\frac{\epsilon^{(0)}_{B,W}(n)}{H_i\alpha_i}\sum_{p\in\mathcal{P}_i}\frac{1}{p^2}\sum_q\frac{|b_{j(q)}|}{q^2},\\
&\frac{\epsilon^{(0)}_{B,W}(n)|b_i|}{H_i\alpha_i}\sum_{p\in\mathcal{P}_i}\frac{1}{p^2}.
\end{aligned}$$

Each is $o_{n;B,W,\mathcal{P}}(1)$ uniformly in i, by (8.68), $\sum_q |b_{j(q)}|/q \ll 1$, and the fact that every positive-cell α_i is bounded below by a fixed multiple of δ, whereas the low-cell value is α_0. We also keep the literal endpoints separate. For $\|b\|_\sharp \le 1$, the endpoint ledger (7.51) gives, orientation by orientation,

$$\begin{aligned}
\frac{|E_i^{JI,\mathrm{end}}|}{\alpha_i} &\ll_{B,W} \frac{L}{n}\frac{\#\mathcal{P}_i}{H_i\alpha_i}\sum_q|b_{j(q)}| \ll_{B,W,\delta,\eta} \frac{y^2L}{n},\\
\frac{|E_i^{IJ,\mathrm{end}}|}{\alpha_i} &\ll_{B,W} \frac{L}{n}\frac{\#\mathcal{P}_i}{H_i\alpha_i}\sum_q|b_{j(q)}| \ll_{B,W,\delta,\eta} \frac{y^2L}{n},\\
\frac{|E_i^{JJ,\mathrm{end}}|}{\alpha_i} &\ll_{B,W} \frac{L^2}{n}\frac{\#\mathcal{P}_i}{H_i\alpha_i}\sum_q|b_{j(q)}| \ll_{B,W,\delta,\eta} \frac{y^2L^2}{n},\\
\frac{|E_i^{\mathrm{diag,end}}|}{\alpha_i} &\ll_{B,W} \frac{L^2}{n}\frac{\#\mathcal{P}_i}{H_i}\frac{|b_i|}{\alpha_i} \ll_{B,W,\delta,\eta} \frac{yL^2}{n}.
\end{aligned}$$

Indeed $\#\mathcal{P}_i \le y$, $\sum_q |b_{j(q)}| \le y$, and $|b_i| \le \alpha_i$. Moreover, $H_i\alpha_i = \sum_{p\in\mathcal{P}_i} t_p/p$ tends to δ on the low cell and to $|I_i|$ on each positive cell, while H_i is bounded below on every permitted cell. Hence the four displays imply the common relative bound

$$\max_i \frac{|E_i^{\mathrm{endpoint}}|}{\alpha_i} \ll_{B,W,\mathcal{P}} \frac{y^2L^2}{n} = o_{n;B,W,\mathcal{P}}(1).$$

Thus all box dependence in the four directions is contained in one remainder $\epsilon^{\sharp}_{B,W,\mathcal{P}}(n) \to 0$ for each fixed mesh, and each direction separately is relative at the moving low row. Consequently

$$|E_i^{\mathrm{pow}}| \le \alpha_i \left\{ \frac{C^{\sharp}_{\mathrm{pow}}}{W} + \epsilon^{\sharp}_{B,W,\mathcal{P}}(n) \right\}, \tag{8.69}$$

where $C^{\sharp}_{\mathrm{pow}}$ is independent of the compact tilt box and of the subsequently chosen fine mesh. Finally (8.48), the marked head/physical rows, and the exact removal in (8.50) give, for $\|b\|_\sharp \le 1$,

$$|(U_\xi)_{i\nu}| \ll_{B,W} \frac{H_i}{L}, \qquad \sum_j H_j |b_j| \le \sum_j H_j \alpha_j \ll 1, \qquad \|\Gamma_\xi^{-1}\| \ll_W 1.$$

Consequently

$$|E_i^{\mathrm{head/phys}}| \ll_{B,W} \frac{1}{L^2} \le \alpha_i \, O_{B,W}\!\left(\frac{\log L}{\delta L^2}\right) = o_{n;B,W}(\alpha_i).$$

For reference, the relative estimates just proved can be read in the following ledger; the displayed quantity is always the supremum over $\|b\|_\sharp \le 1$:

source	$\max_i \lvert (E_n b)_i\rvert/\alpha_i$
leading h, K quadrature	$o_{n;W,\mathcal{P}}(1)$
squarefree marked off-diagonal	$O_{B,W}(\log L/(\delta L))$
guard deletion	$o_{n;B,W,\mathcal{P}}(1)$
squarefree Bernoulli diagonal	$O(1/W)$
prime powers	$C^{\sharp}_{\mathrm{pow}}/W + \epsilon^{\sharp}_{B,W,\mathcal{P}}(n)$
head/physical Schur block	$O_{B,W}(\log L/(\delta L^2))$
arithmetic/continuum projection	$o_{n;W,\mathcal{P}}(1)$.

Thus every vanishing row error is $o(\alpha_i)$, including the low row $i = 0$; no additive $o(1)$ is divided by α_0 after the fact. It follows that the sharp discrete operator is a perturbation, in the sharp norm, whose nonvanishing part is $(C_{\mathrm{sf}} + C^{\sharp}_{\mathrm{pow}})/W$, with every dependence on the later compact box and fixed mesh in a remainder tending to zero with n. Equivalently, before either gauge projection is applied,

$$\left\| S_n^{-1} \mathsf{B}^Z_\xi S_n - S_c^{-1} B_{t_0,\mathcal{P}} S_c \right\|_{\infty\to\infty} \le \frac{C_{\mathrm{sf}} + C^{\sharp}_{\mathrm{pow}}}{W} + o_{n;B,W,\mathcal{P}}(1). \tag{8.70}$$

Step 4: identification of the varying sharp gauges and transfer of the inverse. We finish by identifying the two quotient spaces explicitly. Put

$$\Pi_n := S_n^{-1} P_{\alpha,n} S_n, \qquad \mathcal{R}_n := \operatorname{Ran} \Pi_n,$$
$$\mathcal{T}_{n,\xi} := \Pi_n S_n^{-1} \mathsf{B}^Z_\xi S_n \Pi_n, \qquad \mathcal{T}_c := \Pi_c S_c^{-1} B_{t_0,\mathcal{P}} S_c \Pi_c.$$

Equations (8.41) and (8.70), together with the uniform boundedness of the continuum operator, give the explicit projected comparison

$$\|\mathcal{T}_{n,\xi} - \mathcal{T}_c\|_{\infty\to\infty} \le \frac{C_{\mathrm{sf}} + C^{\sharp}_{\mathrm{pow}}}{W} + o_{n;B,W,\mathcal{P}}(1). \tag{8.71}$$

All operators in this display are extended to the full coordinate space by the displayed projections. Put

$$\varepsilon_n := \|\Pi_n - \Pi_c\|_{\infty\to\infty} = o_n(1).$$

The exact formulas preceding (8.41) show that both projections have kernel $\operatorname{span}\{\mathbf{1}\}$ and hence rank s. For $x \in \mathcal{R}_c$, the restriction

$$U_n := \Pi_n|_{\mathcal{R}_c} : \mathcal{R}_c \longrightarrow \mathcal{R}_n$$

satisfies

$$\|U_n x - x\|_\infty = \|(\Pi_n - \Pi_c)x\|_\infty \le \varepsilon_n \|x\|_\infty.$$

Thus U_n is injective when $\varepsilon_n < 1$, and equality of the two ranks makes it onto. Moreover,

$$\|U_n\| \le 1 + \varepsilon_n, \qquad \|U_n^{-1}\| \le (1-\varepsilon_n)^{-1}. \tag{8.72}$$

If $y = U_n x \in \mathcal{R}_n$, then, under the ambient coordinate identification,

$$\|U_n^{-1} y - y\|_\infty = \|x - \Pi_n x\|_\infty \le \frac{\varepsilon_n}{1-\varepsilon_n}\|y\|_\infty.$$

For $x \in \mathcal{R}_c$, the exact expansion inside U_n^{-1} is

$$\mathcal{T}_{n,\xi} U_n x - U_n \mathcal{T}_c x = (\mathcal{T}_{n,\xi} - \mathcal{T}_c) U_n x + \mathcal{T}_c (U_n - I)x + (I - \Pi_n)\mathcal{T}_c x.$$

The last two terms are each at most $\varepsilon_n \|\mathcal{T}_c\|\,\|x\|_\infty$, up to the harmless factor $1+\varepsilon_n$ in the first one. Equations (8.72) and (8.71) therefore give the pulled-back comparison

$$\left\| U_n^{-1}(\mathcal{T}_{n,\xi}|_{\mathcal{R}_n}) U_n - \mathcal{T}_c|_{\mathcal{R}_c} \right\|_{\infty\to\infty} \le \frac{C_{\mathrm{sf}} + C_{\mathrm{pow}}^{\sharp}}{W} + o_{n;B,W,\mathcal{P}}(1).$$

We enlarge the structural threshold W_0 once and for all, before the actual W is selected, so that the box-independent $1/W$ term is smaller than $(4C_\sharp^{(0)})^{-1}$ for every $W \ge W_0$. Only after this choice are the compact box and mesh fixed and n taken large enough to absorb the remainder. The Neumann series and (8.62) now give

$$\left\| (\mathcal{T}_{n,\xi}|_{\mathcal{R}_n})^{-1} \right\|_{\infty\to\infty} \le 2C_\sharp^{(0)} + o(1).$$

Finally put $x = S_n^{-1} q$ and $v = S_n^{-1} u$. Because $q, u \in \mathcal{G}_n$, the exact projection formula gives $x, v \in \mathcal{R}_n$, and (8.58) becomes $\mathcal{T}_{n,\xi} x = v$. The hypothesis gives $\|v\|_\infty \le Cw$, while $\|x\|_\infty = \|q\|_\sharp$. Applying the preceding inverse bound therefore yields $\|q\|_\sharp \le C_\sharp w$, which is (8.59). This also makes the order of the choices W, box, mesh, and n explicit. □

Lemma 8.6 (Augmented weighted compensated-score lemma)**.** *Uniformly on every fixed compact effective-tilt box, one has*

$$\sum_{W<p\le y} \frac{|g_p|}{p} \ll w, \qquad V \asymp w^2. \tag{8.73}$$

Let Z and $\mathcal{C}_\xi^Z$ be as in (8.23). First Schur-project the raw score $S_g = \sum_p g_p v_p$ off Z, and then regress it on the gauge-fixed quotient band space $\mathcal{G}$. Explicitly, put

$$b_i^Z := \mathrm{Cov}_\xi(\Omega_i, S_g) - \mathrm{Cov}_\xi(\Omega_i, Z)\Gamma_\xi^{-1}\,\mathrm{Cov}_\xi(Z, S_g),$$

$$P_{\alpha,n} D^{-1} \frac{\mathcal{C}_\xi^Z}{q_n} P_{\alpha,n} q^{\mathrm{reg}} = P_{\alpha,n} D^{-1} b^Z, \qquad q^{\mathrm{reg}} \in \mathcal{G}, \tag{8.74}$$

$$a_Z := \Gamma_\xi^{-1}\,\mathrm{Cov}_\xi\left(Z, S_g - \sum_j q_j^{\mathrm{reg}} \Omega_j \right),$$

$$C_g := S_g - \sum_j q_j^{\mathrm{reg}} \Omega_j - a_Z^T Z, \qquad c_p := g_p - q_{j(p)}^{\mathrm{reg}}. \tag{8.75}$$

Thus the Schur operation, quotient regression, and gauge occur in the displayed order. The number c_p is the valuation coefficient before the explicit finite nuisance regression $a_Z^T Z$; that nuisance block is kept separate below. In the display below, the vertical bar denotes the Schur-complement *variance after orthogonal regression on the named nuisance scores; it is not conditioning on a probability-zero event. Then*

$$\|c\|_\infty \le C_{\mathrm{cmp}} w, \qquad \sum_p \frac{|c_p|}{p} \le C_{\mathrm{cmp}} w, \quad \sum_p \frac{c_p^2}{p} \le C_{\mathrm{cmp}} w^2, \tag{8.76}$$

$$\mathrm{Var}_{\nu_{n,\xi}}\left(\sum_p c_p v_p(m) \,\middle|\, \mathcal{G}, Z\right) \asymp_W w^2, \tag{8.77}$$

$$\left|\mathrm{Cov}_{\nu_{n,\xi}}\left(v_p(m), C_g(m)\right)\right| \ll_W \frac{w}{p} \qquad (W < p \le y). \tag{8.78}$$

The contribution of prime powers to (8.77) *is*

$$\frac{C_{\mathrm{rel}}(C_{\mathrm{cmp}})}{W} w^2 + \epsilon^{\mathrm{rel}}_{B,W,C_{\mathrm{cmp}}}(n) w^2, \tag{8.79}$$

rather than an additive error independent of w. Here C_{cmp} is the continuum regression constant, fixed before W and independent of the compact tilt radius; all finite-head and box dependence is in thresholds or in the displayed vanishing remainder, which is defined explicitly in (7.59)*.*

Proof. Mertens' theorem and (8.20) give the first assertion with uniform constants. In the continuum low cell,

$$\int_{t_0}^{\delta} (t - \alpha_0^{(c)})^2 \frac{dt}{t} = \frac{\delta^2 - t_0^2}{2} - \frac{(\delta - t_0)^2}{\log(\delta/t_0)} = \frac{\delta^2}{2} + o(1). \tag{8.80}$$

For a positive cell $I = [a, a+\ell]$, direct integration and the two-sided regular-mesh bounds give

$$\int_I (t - \alpha_I)^2 \frac{dt}{t} \asymp \frac{\ell^3}{a}, \qquad \int_I |t - \alpha_I| \frac{dt}{t} \asymp \frac{\ell^2}{a}. \tag{8.81}$$

The merged terminal cell satisfies the same bounds. Summing the geometric mesh and applying Mertens gives

$$\sum_p |g_p|/p \ll \delta + \eta, \qquad V \asymp \delta^2 + \eta^2 \asymp w^2.$$

We next justify the sharp right-side estimate for the *non-step* coefficient $g_p = \alpha_{j(p)} - t_p$. This does not follow merely by applying the preceding step-input operator estimate, since on the moving low cell one can have $|g_p| \gg \alpha_0$. All prime sums in this paragraph range over $W < p \le y$. Put

$$B_1(g) := \sum_p \frac{|g_p|}{p}, \qquad G_1(g) := \sum_p \frac{t_p|g_p|}{p}, \qquad B_2(g) := \sum_p \frac{|g_p|}{p^2}, \qquad G_2(g) := \sum_p \frac{t_p|g_p|}{p^2}.$$

The cell moment calculations in (8.80)– (8.81), followed by weighted Mertens summation, give

$$\begin{aligned} &\|g\|_\infty \ll w, \qquad B_1(g) \ll w, \qquad G_1(g) \ll w, \\ &G_2(g) \le W^{-1} G_1(g) \ll w/W, \end{aligned} \tag{8.82}$$

$$\frac{1}{H_i} \sum_{p \in \mathcal{P}_i} \frac{t_p|g_p|}{p} \ll w\alpha_i \qquad (0 \le i \le s). \tag{8.83}$$

For the two unweighted reciprocal-square quantities needed below one has, for the fixed regular mesh,

$$B_2(g) = O_W(\alpha_0) + o_n(\alpha_0), \tag{8.84}$$

$$d_i(g) := \frac{1}{H_i} \sum_{p\in\mathcal{P}_i} \frac{|g_p|}{p^2} = o_{n;W,\delta,\eta,\mathcal{P}}(w\alpha_i) \quad \text{uniformly in } i. \tag{8.85}$$

Here is the low-row calculation, where these assertions are not a consequence of a step bound. Since $|g_p| \le \alpha_0 + t_p$ on $\mathcal{P}_0$,

$$d_0(g) \le \frac{\alpha_0}{H_0} \sum_{p>W} \frac{1}{p^2} + \frac{1}{H_0 \log y} \sum_{p>W} \frac{\log p}{p^2} = o_n(w\alpha_0).$$

Indeed, after division by $w\alpha_0$, the two terms are respectively $O_W((wH_0)^{-1})$ and $O_W((w\delta \log y)^{-1})$. The same estimate before division by H_0 proves the low-cell contribution to (8.84). Every positive cell starts at y^δ, so its contribution to either display is exponentially small relative to α_0. This proves (8.84)– (8.85) without losing a factor α_0^{-1}.

We now contract every covariance term against g. The leading squarefree band output is

$$R_i^{h,K}(g) := \frac{1}{H_i} \left\{ \sum_{p\in\mathcal{P}_i} \frac{h(t_p)g_p}{p} + \sum_{p\in\mathcal{P}_i} \sum_{q\neq p} \frac{K(t_p,t_q)g_q}{pq} \right\}.$$

Exact harmonic centering gives $\sum_{p\in\mathcal{P}_i} g_p/p = 0$. Since $h(t) - 1 = O(t)$, $|K(s,t)| \ll st$, and $H_i^{-1} \sum_{p\in\mathcal{P}_i} t_p/p = \alpha_i$, (8.82)– (8.83) give

$$\left| R_i^{h,K}(g) \right| \ll \frac{1}{H_i} \sum_{p\in\mathcal{P}_i} \frac{t_p|g_p|}{p} + \alpha_i G_1(g) \ll w\alpha_i. \tag{8.86}$$

In particular this includes the low-row estimate that was formerly implicit. The squarefree marked error (8.63) contracts to

$$|R_i^{\mathrm{sf,err}}(g)| \ll_{B,W} \frac{B_1(g)}{L} \ll_{B,W} \frac{w}{L}, \qquad \max_i \frac{|R_i^{\mathrm{sf,err}}(g)|}{w\alpha_i} \ll_{B,W} \frac{\log L}{\delta L} = o(1).$$

The Bernoulli diagonal is bounded by $d_i(g)$, hence is $o(w\alpha_i)$ by (8.85). Guard deletion is also relative: apply (7.63) to $g/\|g\|_\infty$ and use $\|g\|_\infty \ll w$. Its output is therefore $o(w\alpha_0) = o(w\alpha_i)$ on every row.

For completeness, decompose $V_p = I_p + J_p$ and use (7.34)– (7.36). The box-independent main terms in the four orientations satisfy the following bounds, where $R_i^{JI}, R_i^{IJ}, R_i^{JJ}, R_i^{\mathrm{diag}}$ denote their respective contributions to $(D^{-1}\operatorname{Cov}_\xi(\Omega, S_g))_i$:

$$\begin{aligned}
|R_i^{JI}(g)|_{\mathrm{main}} &\ll \frac{1}{H_i} \sum_{p\in\mathcal{P}_i} \frac{t_p}{p^2} G_1(g) \ll \frac{w\alpha_i}{W}, \\
|R_i^{IJ}(g)|_{\mathrm{main}} &\ll \frac{1}{H_i} \sum_{p\in\mathcal{P}_i} \frac{t_p}{p} G_2(g) \ll \frac{w\alpha_i}{W}, \\
|R_i^{JJ}(g)|_{\mathrm{main}} &\ll \frac{1}{H_i} \sum_{p\in\mathcal{P}_i} \frac{t_p}{p^2} G_2(g) \ll \frac{w\alpha_i}{W^2}, \\
|R_i^{\mathrm{diag}}(g)|_{\mathrm{main}} &\ll d_i(g) = o_n(w\alpha_i).
\end{aligned}$$

Let $\epsilon_0 = \epsilon_{B,W}^{(0)}(n)$ be the pair remainder in (7.52). After division by $w\alpha_i$, its four contractions are bounded, respectively, by

$$\epsilon_0 \frac{1}{H_i\alpha_i} \sum_{p\in\mathcal{P}_i} \frac{1}{p^2} \frac{B_1(g)}{w}, \quad \epsilon_0 \frac{B_2(g)}{w\alpha_i}, \quad \epsilon_0 \frac{1}{H_i\alpha_i} \sum_{p\in\mathcal{P}_i} \frac{1}{p^2} \frac{B_2(g)}{w}, \quad \epsilon_0 \frac{d_i(g)}{w\alpha_i}.$$

The first expression tends to zero by (8.68) and (8.82); the third and fourth do so by (8.84)–(8.85). For the second expression, the low row is $O_W(\epsilon_0/w)$, while on every positive row $\alpha_i \gg \delta$; it too tends to zero after the already chosen positive w, W, box, and mesh have been fixed. This is threshold dependence only and creates no box-dependent main constant. Finally, the literal endpoint terms obey

$$\max_i \frac{|R_i^{\mathrm{endpoint}}(g)|}{w\alpha_i} \ll_{B,W,\delta,\eta,\mathcal{P}} \frac{y^2L^2}{n} = o(1).$$

Indeed $\#\mathcal{P}_i \le y$, $\sum_q |g_q| \ll wy$, and $H_i\alpha_i$ tends to δ on the low cell and is bounded below on every fixed positive cell. Thus all JI, IJ, JJ, diagonal, tail, endpoint, squarefree, and guard terms have now been contracted against the actual non-step score.

The finite head and physical Schur subtraction is relative as well. We record its raw nuisance vector here so it is not inferred from the later post-regression calculation. Before guard deletion, (7.64) and centered component-weight cancellation give

$$|\operatorname{Cov}_\xi(H_\nu^\circ, S_g)| \ll_{B,W} \frac{1}{L}\sum_p \frac{|g_p|}{p} \ll_{B,W} \frac{w}{L}.$$

The physical coordinate has the same bound by (7.61). Since $\|S_g\|_\infty \ll_W wL$, (7.65) and (7.67) show that guard deletion changes either covariance by $O_{B,W}(wyL^{A+1}/n) = o_{n;B,W}(w/L)$. Hence

$$\|\operatorname{Cov}_\xi(Z, S_g)\| \ll_{B,W} \frac{w}{L} + o_{n;B,W}(w/L).$$

Together with $\|(U_\xi)_{i,*}\| \ll_{B,W} H_i/L$, (8.48), and $\|\Gamma_\xi^{-1}\| \ll_W 1$, this gives

$$\frac{1}{H_i}\left|\left(U_\xi\Gamma_\xi^{-1}\operatorname{Cov}_\xi(Z,S_g)\right)_i\right| \ll_{B,W} \frac{w}{L^2} + o_n(w/L^2) = o_{n;B,W}(w\alpha_i). \tag{8.87}$$

Combining (8.86)– (8.87), we have proved the explicit sharp ledger

$$\max_i \frac{|(D^{-1}b^Z)_i|}{w\alpha_i} \le C_{\mathrm{rhs}} + \frac{C_{\mathrm{pow}}^{\mathrm{raw}}}{W} + r_{n;B,W,\mathcal{P}}, \qquad r_{n;B,W,\mathcal{P}} \longrightarrow 0.$$

Here $C_{\mathrm{rhs}}, C_{\mathrm{pow}}^{\mathrm{raw}}$ depend only on the structural Dickman and regular-mesh bounds; in particular they are independent of the later box radius and of the permitted mesh. All dependence on the fixed box and mesh is confined to the vanishing remainder and its n-threshold. Moreover, (8.41) shows that $S_n^{-1}P_{\alpha,n}S_n$ is uniformly bounded, so projecting the right side in (8.74) preserves this sharp estimate. Increase only the corresponding n-threshold until $r_{n;B,W,\mathcal{P}} \le 1$. The right-side constant passed to the weighted inverse can therefore be chosen from C_{rhs}, the already fixed structural $1/W_0$ bound, and one spare unit, independently of the actual W, box radius, or permitted mesh. Applying Lemma 8.5 gives

$$|q_0^{\mathrm{reg}}| \le C_\sharp w\alpha_0, \qquad |q_j^{\mathrm{reg}}| \le C_\sharp w\alpha_j \quad (j \ge 1). \tag{8.88}$$

We next prove, rather than assume, the sharper nuisance estimate needed for the lower Schur bound. Put

$$P_{\mathrm{med}} := S_g - q^{\mathrm{reg}}\cdot\Omega = \sum_p c_pV_p, \qquad c_p = g_p - q_{j(p)}^{\mathrm{reg}}.$$

Equations (8.73) and (8.88), together with $\sum_j H_j\alpha_j = O(1)$, give

$$\|c\|_\infty \ll w, \qquad \sum_p \frac{|c_p|}{p} \le \sum_p \frac{|g_p|}{p} + \sum_j H_j|q_j^{\mathrm{reg}}| \ll w.$$

Moreover, the quadratic coefficient bound used below is the direct calculation

$$\sum_p \frac{c_p^2}{p} \le 2V + 2\sum_j H_j(q_j^{\mathrm{reg}})^2 \le 2V + 2C_\sharp^2 w^2 \sum_j H_j \alpha_j^2 \ll w^2.$$

First omit the numerical guards and decompose the full tilted mixture into its finitely many exact component cells C_a, with weights λ_a. By (7.64), for a fixed reference component one may write

$$\mathbb{E}_{C_a,\xi} V_p = m_p + e_a(p), \qquad |e_a(p)| \ll_{B,W} \frac{1}{pL},$$

uniformly in a, p. Each head indicator H_ν° is constant on C_a; write that value as $h_{\nu,a}$ and its mixture mean as $\bar{h}_\nu$. The law of total covariance and $\sum_a \lambda_a(h_{\nu,a} - \bar{h}_\nu) = 0$ give the exact cancellation

$$\begin{aligned}\operatorname{Cov}_\xi(H_\nu^\circ, P_{\mathrm{med}}) &= \sum_a \lambda_a(h_{\nu,a} - \bar{h}_\nu) \sum_p c_p\{m_p + e_a(p)\} \\ &= \sum_a \lambda_a(h_{\nu,a} - \bar{h}_\nu) \sum_p c_p e_a(p) \ll_{B,W} \frac{1}{L}\sum_p \frac{|c_p|}{p} \ll_{B,W} \frac{w}{L}. \end{aligned} \tag{8.89}$$

The physical coordinate satisfies the same bound directly by (7.61):

$$|\operatorname{Cov}_\xi(R, P_{\mathrm{med}})| \le \sum_p |c_p|\,|\operatorname{Cov}_\xi(R, V_p)| \ll_{B,W} \frac{w}{L}. \tag{8.90}$$

It remains to check that numerical guards do not destroy these cancellations. Every coordinate of Z is bounded by a constant depending only on W, while

$$\|P_{\mathrm{med}}\|_\infty \le \|c\|_\infty \sum_{W<p\le y} V_p \ll_W wL.$$

Equations (7.65) and (7.67) therefore show that guard deletion changes each covariance in (8.89)–(8.90) by

$$O_{B,W}\left(\frac{wyL^{A+1}}{n}\right) = o_{n;B,W}(w/L).$$

Since $\dim Z$ is fixed after W, we have proved the vector estimate

$$\|\operatorname{Cov}_\xi(Z, P_{\mathrm{med}})\| \ll_{B,W} \frac{w}{L} + o_{n;B,W}(w/L). \tag{8.91}$$

Together with $\|\Gamma_\xi^{-1}\| \ll_W 1$, this gives

$$\|a_Z\| = \|\Gamma_\xi^{-1}\operatorname{Cov}_\xi(Z, P_{\mathrm{med}})\| \ll_{B,W} \frac{w}{L} + o_n(w/L) \le C_W w$$

after increasing only the n-threshold for the already fixed box. Thus the nonvanishing comparison constant may still be chosen independently of the box radius. Equations (8.73) and (8.88) imply all three bounds in (8.76). The constants used for the medium-prime vector are the Mertens main-term constants and the continuum weighted inverse underlying $C_\sharp$. The preceding relative display shows that finite nuisance regression changes the medium vector only by a vanishing relative term; its nonvanishing finite-dimensional part is recorded separately in a_Z. Hence C_{cmp} is fixed from the continuum regression before W, the tilt-box radius, or the final mesh is chosen.

For the lower Schur bound we again keep the continuum and arithmetic objects separate. Define

$$V_c := \int_{t_0}^1 |\alpha_{\mathcal{P}}^{(c)}(t) - t|^2 \,\frac{dt}{t}, \qquad \alpha_{\mathcal{P}}^{(c)}(t) := \alpha_j^{(c)} \quad (t \in I_j).$$

The moment calculations above and weighted Mertens summation give

$$V = V_c + o_{n;W,\mathcal{P}}(w^2), \qquad V_c \asymp w^2. \tag{8.92}$$

To transfer the actual fitted vector $q^{\mathrm{reg}} \in \mathcal{G}_n$, put

$$x_n := S_n^{-1} q^{\mathrm{reg}}, \qquad x_c := \Pi_c x_n, \qquad q^{(c)} := S_c x_c.$$

Then $q^{(c)} \in \mathcal{G}_c$, and (8.41) together with (8.88) gives

$$\max_j \left| \frac{q_j^{(c)}}{\alpha_j^{(c)}} - \frac{q_j^{\mathrm{reg}}}{\alpha_j} \right| = o_{n;W,\mathcal{P}}(w). \tag{8.93}$$

Let $q_{\mathcal{P}}^{(c)}(t) = q_j^{(c)}$ on I_j. Harmonic centering of $t - \alpha_j^{(c)}$ on each cell gives the exact continuum identity

$$\begin{aligned} &\inf_{\lambda \in \mathbb{R}} \int_{t_0}^1 \left| \alpha_{\mathcal{P}}^{(c)}(t) - q_{\mathcal{P}}^{(c)}(t) - \lambda t \right|^2 \frac{dt}{t} \\ &\quad = \inf_{\lambda \in \mathbb{R}} \left\{ \|(1-\lambda)\alpha^{(c)} - q^{(c)}\|_{D_c}^2 + \lambda^2 V_c \right\} \geq \frac{A_c V_c}{A_c + V_c} \gg V_c, \end{aligned} \tag{8.94}$$

where

$$A_c := (\alpha^{(c)})^T D_c \alpha^{(c)} \asymp 1, \qquad (\alpha^{(c)})^T D_c q^{(c)} = 0.$$

The second relation is the exact continuum gauge. Thus no arithmetic orthogonality is inserted into a continuum integral.

The exact finite-n identity

$$\sum_{W < p \leq y} t_p v_p(m) = \frac{L + R(m) - \sum_{p \leq W} v_p(m) \log p}{\log y} \tag{8.95}$$

shows that, after row centering, the physical and head columns subtract the function t exactly at finite n. Hence the distance from the kernel direction in (8.94) is the one relevant after the joint nuisance Schur complement.

We spell out the discrete-to-continuum quadratic transfer for the actual non-step coefficient. On $I_j \cap [t_0, 1]$, put

$$c_{\mathcal{P}}^{(c)}(t) := \alpha_j^{(c)} - t - q_j^{(c)},$$

and extend its value at t_0 linearly to $(0, t_0)$, exactly as in (8.30); denote the resulting function on $(0,1]$ by $\tilde{c}_{\mathcal{P}}^{(c)}$. Cellwise Mertens quadrature, (8.92), and (8.93) first give

$$\sum_{W < p \leq y} \frac{|(\alpha_{j(p)} - t_p - q_{j(p)}^{\mathrm{reg}}) - c_{\mathcal{P}}^{(c)}(t_p)|}{p} = o_{n;W,\mathcal{P}}(w),$$

$$\sum_{W < p \leq y} \frac{|(\alpha_{j(p)} - t_p - q_{j(p)}^{\mathrm{reg}}) - c_{\mathcal{P}}^{(c)}(t_p)|^2}{p} = o_{n;W,\mathcal{P}}(w^2).$$

Here W, δ, η, and hence $w > 0$, have already been fixed before $n \to \infty$. Thus the relative center convergence in (8.40) is legitimately $o_n(w)$ at this stage; no estimate uniform as $w \downarrow 0$ is being asserted. The first line controls the marked off-diagonal error, and the second, together with boundedness of the continuum covariance operator and Cauchy–Schwarz, changes its quadratic form by $o(w^2)$. The two-index marked formula then gives, uniformly on each fixed box,

$$\mathrm{Var}_\xi \left(\sum_{W < p \leq y} c_p I_p \right) = \langle \tilde{c}_{\mathcal{P}}^{(c)}, A \tilde{c}_{\mathcal{P}}^{(c)} \rangle_{\mathcal{H}} + O_{B,W}(w^2/L) + O(w^2/W) + o_{n;B,W,\mathcal{P}}(w^2).$$

Here the $O(w^2/W)$ term is precisely the squarefree diagonal correction; it is retained rather than hidden in the little-oh term. At the low boundary,

$$|c_{\mathcal{P}}^{(c)}(t_0)| \le \alpha_0^{(c)} + t_0 + |q_0^{(c)}| = O(\alpha_0^{(c)} + t_0) = o_n(w).$$

Consequently the low extension changes both the squared norm and the quadratic form by $o_n(w^2)$, by the calculation in (8.32). Therefore (8.12), (8.94), and (8.92) give a fixed positive multiple of w^2 before the following relative perturbations are subtracted. The finite nuisance Schur loss is, by (8.91) and (8.48),

$$\mathrm{Cov}_\xi(Z, S_g - q^{\mathrm{reg}} \cdot \Omega)^T \, \Gamma_\xi^{-1} \, \mathrm{Cov}_\xi(Z, S_g - q^{\mathrm{reg}} \cdot \Omega) = O_{B,W}(w^2/L^2) + o_{n;B,W,\mathcal{P}}(w^2).$$

Thus the continuum Poincaré gap is transferred at the relative w^2 scale before the prime-power replacement is made. The upper bound in (8.77) follows directly from the unprojected score and (8.73).

At the w^2 scale every transfer is relative. From the marked squarefree errors and the coefficient bounds already proved,

$$\begin{aligned} E_{\mathrm{sqfree}} &\ll_{B,W} \frac{1}{L} \left\{ \left(\sum_p \frac{|c_p|}{p} \right)^2 + \sum_p \frac{c_p^2}{p} \right\} = O_{B,W}(w^2/L), \\ E_{\mathrm{sf,diag}} &\ll \sum_p \frac{c_p^2}{p^2} \le \frac{1}{W} \sum_p \frac{c_p^2}{p} = O(w^2/W), \\ E_{\mathrm{disc}} &\le \varepsilon_{W,\delta,\eta}(n) \left\{ \left(\sum_p \frac{|c_p|}{p} \right)^2 + \sum_p \frac{c_p^2}{p} \right\} = o_n(w^2), \end{aligned} \tag{8.96}$$

where quantitative Mertens summation gives $\varepsilon_{W,\delta,\eta}(n) \ll 1/\log L + \log L/L + e^{-c\sqrt{L}}$. The endpoint terms in the four-mark formula are included in this ε; the cofactor margin $n/y^4 = n^{1/9}$ makes their total $o_n(w^2)$ for the fixed mesh. For the nuisance block, (8.91), the gap $\Gamma_\xi^{-1} = O_W(1)$, and the exact physical removal in (8.50) give

$$E_{\mathrm{head/phys}} = o_n(w^2). \tag{8.97}$$

These displays include respectively the off-diagonal squarefree, squarefree-diagonal, discretization, and head/physical errors; none is being replaced by a generic additive $o_n(1)$.

It remains to justify the relative prime-power statement. Write $I_p = 1_{p|m}$ and $J_p = v_p - I_p$. The product-weighted aggregate estimate (7.34), with its transposed and two-extra-power analogues, gives, for $p \ne q$,

$$|\mathrm{Cov}(J_p, I_q)| \le C_{\mathrm{pow}} \frac{t_p t_q}{p^2 q} + \frac{\epsilon_{B,W}(n)}{p^2 q} + O_{B,W}(L/n),$$

with product-weighted analogues for the diagonal, transpose, and $\mathrm{Cov}(J_p, J_q)$. Hence

$$\begin{aligned} &\left| \mathrm{Var}\left(\sum_p c_p v_p \right) - \mathrm{Var}\left(\sum_p c_p I_p \right) \right| \\ &\qquad \le \frac{C_{\mathrm{rel}}(C_{\mathrm{cmp}})}{W} w^2 + \epsilon^{\mathrm{rel}}_{B,W,C_{\mathrm{cmp}}}(n) w^2. \end{aligned}$$

This is (7.60), applied with (8.76). Together with the $O(w^2/W)$ squarefree diagonal in (8.96) and (8.97), this proves (8.79). For clarity, the complete relative-error ledger is therefore

contribution	bound
$E_{\mathrm{sf,off}}$	$O_W(w^2/L)$
$E_{\mathrm{sf,diag}}$	$O(w^2/W)$
E_{pow}	$C_{\mathrm{rel}}(C_{\mathrm{cmp}})w^2/W + \epsilon^{\mathrm{rel}}_{B,W,C_{\mathrm{cmp}}}(n)w^2$
E_{disc}	$o_{n;W,\mathcal{P}}(w^2)$
E_{head}	$o_{n;B,W,\mathcal{P}}(w^2)$
E_{phys}	$o_{n;B,W,\mathcal{P}}(w^2)$.

The last two entries are listed separately: the head term follows from the finite-dimensional gap (8.48) and the marked head rows, while the physical leading term is removed exactly by (8.50) and its remaining Stieltjes error is $o_n(w^2)$. The constant C_{cmp}, and hence $C_{\mathrm{rel}}(C_{\mathrm{cmp}})$, was fixed from the continuum regression before W, the mesh, or the tilt box; in particular it is independent of all three. Choose W using the box-independent constant $C_{\mathrm{rel}}(C_{\mathrm{cmp}})$, and only after the compact box is known take n large enough to absorb $\epsilon^{\mathrm{rel}}_{B,W,C_{\mathrm{cmp}}}(n)$; the lower Schur bound is then preserved.

Finally we verify the marked row without suppressing an orientation, diagonal, or endpoint term. Fix $W < p \le y$; all covariances in the following calculation are under $\nu_{n,\xi}$. Since $C_g = \sum_q c_q V_q - a_Z^T Z$, first consider the prime part. The squarefree diagonal is $O(|c_p|/p) = O(w/p)$. For $q \ne p$, (8.35) and (8.36) give

$$\sum_{q\ne p} |c_q|\,|\operatorname{Cov}(I_p, I_q)| \ll \frac{t_p}{p}\sum_q \frac{t_q|c_q|}{q} + \frac{\|c\|_\infty}{p}\left\{\frac{C_{\mathrm{sf}}}{W} + \epsilon^{\mathrm{sf}}_{B,W}(n)\right\} \ll_{B,W} \frac{w}{p}. \tag{8.98}$$

Here and below the threshold for the vanishing remainder may depend on the already fixed box. The coefficient estimates (8.76) also imply

$$\sum_q |c_q| \ll wy, \qquad \sum_q \frac{t_q|c_q|}{q} \ll w, \qquad \sum_q \frac{t_q|c_q|}{q^2} \ll \frac{w}{W}.$$

Using $V = I + J$, the four estimates (7.34)–(7.36) now give the following separate contractions:

$$\begin{aligned}
\sum_{q\ne p} |c_q|\,|\operatorname{Cov}(J_p, I_q)| &\ll \frac{w}{p^2} + \epsilon_{B,W}(n)\frac{w}{p^2} + \frac{wyL}{n},\\
\sum_{q\ne p} |c_q|\,|\operatorname{Cov}(I_p, J_q)| &\ll \frac{w}{pW} + \epsilon_{B,W}(n)\frac{w}{pW} + \frac{wyL}{n},\\
\sum_{q\ne p} |c_q|\,|\operatorname{Cov}(J_p, J_q)| &\ll \frac{w}{p^2W} + \epsilon_{B,W}(n)\frac{w}{p^2W} + \frac{wyL^2}{n},\\
|c_p|\,|\operatorname{Var}(V_p) - \operatorname{Var}(I_p)| &\ll \frac{w}{p^2} + \epsilon_{B,W}(n)\frac{w}{p^2} + \frac{wL^2}{n}.
\end{aligned} \tag{8.99}$$

For the last line, $J_p > 0$ implies $I_p = 1$, whence $|\operatorname{Var}(I_p + J_p) - \operatorname{Var}(I_p)| \le 3\mathbb{E}J_p^2$; this is exactly the diagonal estimate in (7.36). All endpoint terms in (8.99) are uniformly $o(w/p)$, since

$$\sup_{p\le y} \frac{pyL^2}{n} \le \frac{y^2L^2}{n} = n^{-5/9}L^2 = o(1).$$

Thus the full prime part of C_g contributes $O_W(w/p)$, uniformly on the fixed box once its n-threshold is chosen. We also spell out the nuisance marked family used here and later. Before guard deletion, (7.64) and the same centered component-weight cancellation as in (8.89) give

$$|\operatorname{Cov}(V_p, H_\nu^\circ)| \ll_{B,W} \frac{1}{pL}$$

for every head coordinate, while (7.61) gives this estimate for the physical coordinate. Guard deletion changes either covariance by at most $O_{B,W}(\delta_\xi L)$, by (7.67). Uniformly for $p \le y$,

$$\delta_\xi L\,(pL) \ll_{B,W} \frac{y^2 L^{A+2}}{n} = o(1),$$

so, in the fixed-dimensional Euclidean norm,

$$\|\operatorname{Cov}_\xi(V_p, Z)\| \ll_{B,W} \frac{1}{pL}. \tag{8.100}$$

Since $\|a_Z\| \ll_W w$, it follows that

$$|\operatorname{Cov}(V_p, a_Z^T Z)| \ll_{B,W} \frac{w}{pL} \ll_W \frac{w}{p}$$

after increasing the threshold for the fixed box. Adding this nuisance row to (8.98) and (8.99) proves (8.78). □

8.6 Exact nonlinear fitting

At the physically and head-fitted point constructed above, all frozen charges are included in the residual. Complete rough-signature equations give $r_p = 0$ for $p > y$, and the head fit gives $r_p = 0$ for $p \le W$. The ordinary logarithm was fitted exactly, so

$$0 = \log Y - \log B_0 - \sum_a x_a \log a = \sum_p r_p \log p.$$

Consequently the following compatibility is an exact finite-n identity:

$$\sum_{W<p\le y} t_p r_p = 0. \tag{8.101}$$

Let

$$\Delta_j := \sum_{p\in\mathcal{P}_j} r_p.$$

After absorbing the marked cost (8.8) into the constant, the strict-rate output (6.55) gives

$$|\Delta_j| \ll_W \frac{NH_j}{L}. \tag{8.102}$$

The scalar target requires the sharper estimate below, rather than a generic $O(N/L)$ bound. By (8.101),

$$\sum_j \alpha_j \Delta_j = \sum_{W<p\le y} (\alpha_{j(p)} - t_p) r_p,$$

$$\left|\sum_j \alpha_j \Delta_j\right| \ll_W \frac{N}{L} \sum_{W<p\le y} \frac{|\alpha_{j(p)} - t_p|}{p} \ll_W \frac{Nw}{L}. \tag{8.103}$$

Proposition 8.7 (Physically centered fixed-partition fit)**.** *For the regular partition above, the active weights can be changed while keeping their total mass, every head valuation, and the ordinary logarithm exact, so that*

$$\sum_{p\in\mathcal{P}_j} r_p = 0 \qquad (0 \le j \le s). \tag{8.104}$$

At the resulting point,

$$|r_p| \le C_{1,W} \frac{N}{pL} \qquad (W < p \le y), \tag{8.105}$$

where $C_{1,W}$ is independent of every sufficiently fine regular partition chosen after W. The top component and protected layer in (8.1) remain literally unchanged, all total coordinate weights lie in $[0,1]$, and the integer equation (8.2) is preserved.

Proof. We define one finite-dimensional map and solve all moments simultaneously. For $u \in \mathcal{G}$, put $\mathscr{B}(u) = \sum_j u_j \Omega_j$. Use the fixed statistic $S_g = \sum_p g_p v_p$ and the nuisance vector Z from (8.23). A gauge-fixed parameter is $\xi = (u, a, \lambda) \in \mathcal{G} \times \mathbb{R}^{e+1} \times \mathbb{R}$, with exponential score

$$\mathscr{B}(u) + a^T Z + \lambda S_g.$$

The gauge coordinates are concrete. If $e_0, \ldots, e_s$ are the standard band vectors, use, for $1 \le j \le s$,

$$q^{(j)} := e_j - \frac{H_j \alpha_j}{H_0 \alpha_0} e_0. \tag{8.106}$$

Since $H_0\alpha_0 = \delta + o(1)$, these vectors are well defined and form a basis of $\mathcal{G} = \{u : \alpha^T D u = 0\}$. Primal coordinates mean coefficients in this basis; dual coordinates mean evaluation on the same basis. On the nuisance block we use the ordinary Euclidean norm. The intrinsic quotient-dual norm is specified below and is the one used for all mesh-uniform estimates. This is merely a nonsingular change of coordinates in (8.21), by (8.95). For a statistic F, put

$$m_F(\xi) = \sum_m z_m(\xi) F(m).$$

To avoid concealing a coordinate factor, regard the quotient moment as an element of the finite dual $\mathcal{G}^*$: its value at $v \in \mathcal{G}$ is the moment of $\mathscr{B}(v)$. The normalized moment map is

$$\mathcal{M}(\xi) := \frac{L}{q_n}\Big([v \mapsto m_{\mathscr{B}(v)}(\xi) - m_{\mathscr{B}(v)}(0)], m_Z(\xi) - m_Z(0), m_{S_g}(\xi) - m_{S_g}(0)\Big).$$

Its prescribed endpoint is

$$\tau := \frac{L}{q_n}\Big([v \mapsto v^T\Delta],\ 0,\ \alpha^T\Delta\Big).$$

To specify the quotient norm without a coordinate convention, put

$$d_D := \frac{L}{q_n} D^{-1}\Delta.$$

Then

$$\tau_{\mathcal{G}}(v) = \langle v, d_D\rangle_D, \qquad \|d_D\|_\infty \ll_W 1,$$

by (8.102); the actual gauge representative is $P_{\alpha,n} d_D$. We define $\|\tau_{\mathcal{G}}\|_\infty := \|P_{\alpha,n} d_D\|_\infty$. The projection is uniformly bounded by (8.22). This target is exact: if the Z-moments do not change, then (8.101) says that changing the band moments by Δ changes the S_g-moment by $\alpha^T\Delta$. The envelopes (8.102) and (8.103) give

$$\|\tau_{\mathcal{G}}\|_\infty + \|\tau_Z\| \ll_W 1, \qquad |\tau_g| \ll_W w.$$

We display the full inverse used at an arbitrary point of the path. Let $\mathscr{Q}$ denote the joint vector of the quotient band scores $(\mathscr{B}(q^{(j)}))_{1\le j\le s}$ and Z, in the concrete coordinates (8.106), and write

$$A_\xi := \mathrm{Cov}_\xi(\mathscr{Q}, \mathscr{Q}), \quad b_\xi := \mathrm{Cov}_\xi(\mathscr{Q}, S_g), \quad d_\xi := \mathrm{Var}_\xi(S_g), \quad \sigma_\xi := d_\xi - b_\xi^T A_\xi^{-1} b_\xi.$$

In these primal/dual gauge coordinates, the Jacobian of $\mathcal{M}$ is exactly the covariance block

$$D\mathcal{M}(\xi) = \begin{pmatrix} A_\xi & b_\xi \\ b_\xi^T & d_\xi \end{pmatrix}. \tag{8.107}$$

For a right side $(v_{\mathscr{Q}}, v_g)$, its inverse is not a sequential nonlinear fit but the single block formula

$$\dot\lambda = \sigma_\xi^{-1}\{v_g - b_\xi^T A_\xi^{-1} v_{\mathscr{Q}}\},$$

$$\dot{\mathscr{Q}} = A_\xi^{-1}\{v_{\mathscr{Q}} - b_\xi \dot{\lambda}\}. \tag{8.108}$$

Thus the quotient and nuisance coordinates are continuously refitted whenever the slow coordinate moves.

We now identify both terms involving A_ξ^{-1} with the operators estimated in the preceding lemmas. This is an exact finite-dimensional calculation. Let $\mathbf{Q} : \mathbb{R}^s \to \mathcal{G}$ be the matrix whose columns are the vectors $q^{(j)}$ in (8.106); thus a concrete quotient coordinate x represents the actual band vector $q = \mathbf{Q}x$. Set

$$\begin{aligned} C_\xi &:= \mathrm{Cov}_\xi(\Omega, \Omega), & U_\xi &:= \mathrm{Cov}_\xi(\Omega, Z), & \Gamma_\xi &:= \mathrm{Cov}_\xi(Z, Z), \\ h_\xi &:= \mathrm{Cov}_\xi(\Omega, S_g), & z_\xi &:= \mathrm{Cov}_\xi(Z, S_g), & \overline{C}_\xi &:= C_\xi - U_\xi \Gamma_\xi^{-1} U_\xi^T. \end{aligned}$$

Thus $\overline{C}_\xi = q_n^{-1} \mathcal{C}_\xi^Z$. Literally in the concrete primal/dual coordinates of (8.106),

$$A_\xi = \begin{pmatrix} \mathbf{Q}^T C_\xi \mathbf{Q} & \mathbf{Q}^T U_\xi \\ U_\xi^T \mathbf{Q} & \Gamma_\xi \end{pmatrix}, \qquad b_\xi = \begin{pmatrix} \mathbf{Q}^T h_\xi \\ z_\xi \end{pmatrix}.$$

Let a joint right side be $(r_{\mathcal{G}}, r_Z)$, and choose any $r^\flat \in \mathbb{R}^{s+1}$ representing its quotient functional: $r_{\mathcal{G}}(v) = v^T r^\flat$ for $v \in \mathcal{G}$. Directly eliminating the nuisance variable in $A_\xi(x, a) = (r_{\mathcal{G}}, r_Z)$, with $q = \mathbf{Q}x$, gives the equivalent pair of equations

$$\begin{aligned} a &= \Gamma_\xi^{-1}(r_Z - U_\xi^T q), \\ v^T \overline{C}_\xi q &= v^T (r^\flat - U_\xi \Gamma_\xi^{-1} r_Z) \qquad (v \in \mathcal{G}). \end{aligned} \tag{8.109}$$

Since the Euclidean annihilator of $\mathcal{G}$ is $\mathrm{span}\{D\alpha\}$, the second line is, in turn, equivalent to the intrinsic projected equation

$$P_{\alpha,n} D^{-1} \overline{C}_\xi P_{\alpha,n} q = P_{\alpha,n} D^{-1}(r^\flat - U_\xi \Gamma_\xi^{-1} r_Z), \qquad q \in \mathcal{G}. \tag{8.110}$$

This equivalence uses the actual band values of q; no norm of the coordinate matrix $\mathbf{Q}$ or its inverse is introduced.

Take first $(r^\flat, r_Z) = (h_\xi, z_\xi)$. Equations (8.109)– (8.110) are then exactly (8.74) and (8.75). If $q^{\mathrm{reg}} = \mathbf{Q}x^{\mathrm{reg}}$, they prove, without an implicit coordinate identification,

$$A_\xi^{-1} b_\xi = (x^{\mathrm{reg}}, a_Z), \qquad \sigma_\xi = \mathrm{Var}_\xi(S_g - q^{\mathrm{reg}} \cdot \Omega - a_Z^T Z) = \mathrm{Var}_\xi(C_g). \tag{8.111}$$

The variance identity follows also by expanding the square and using the two normal equations in (8.109); thus it is an ordinary finite Schur complement, not an additional probabilistic assumption.

For the actual fast target, $v_{\mathscr{Q}} = (\tau_{\mathcal{G}}, 0)$, choose

$$r^\flat = \frac{L}{q_n} \Delta = D d_D.$$

Write $A_\xi^{-1} v_{\mathscr{Q}} = (x^{(0)}, a^{(0)})$ and put $u^{(0)} = \mathbf{Q}x^{(0)} \in \mathcal{G}$. The same exact elimination now reads

$$P_{\alpha,n} D^{-1} \overline{C}_\xi P_{\alpha,n} u^{(0)} = P_{\alpha,n} d_D, \qquad a^{(0)} = -\Gamma_\xi^{-1} U_\xi^T u^{(0)}. \tag{8.112}$$

By (8.26) and $\|P_{\alpha,n} d_D\|_\infty \ll_W 1$,

$$\|u^{(0)}\|_\infty \ll_W 1. \tag{8.113}$$

Moreover $|(U_\xi)_{i,*}| \ll_{B,W} H_i/L$, $\sum_i H_i \ll \log L$, and $\|\Gamma_\xi^{-1}\| \ll_W 1$. Therefore

$$\|a^{(0)}\| \ll_{B,W} \frac{\log L}{L} = o_{n;B,W}(1).$$

In particular the fast nuisance coefficient is bounded by an absolute multiple depending only on W, after increasing the threshold for the already fixed box; its bound does not enter the choice of that box.

We may now apply Lemma 8.6 through the exact identification (8.111). It gives

$$\sigma_\xi \asymp_W w^2, \qquad |q_j^{\mathrm{reg}}| \ll w\alpha_j, \qquad \|a_Z\| \ll_W w, \qquad \|c\|_\infty = \max_p |g_p - q_{j(p)}^{\mathrm{reg}}| \ll w. \tag{8.114}$$

The only nontrivial target pairing in the numerator of (8.108) is now an exact dual pairing:

$$\begin{aligned} |b_\xi^T A_\xi^{-1} v_{\mathscr{Q}}| = |(A_\xi^{-1} b_\xi)^T v_{\mathscr{Q}}| &= \left| \frac{L}{q_n} \sum_j q_j^{\mathrm{reg}} \Delta_j \right| \\ &\ll_W w \sum_j H_j \alpha_j \ll_W w. \end{aligned}$$

Here the nuisance target is zero, (8.102) bounds Δ_j, and $\sum_j H_j\alpha_j = \sum_{W<p\le y} t_p/p = O(1)$. Together with $|v_g| \ll_W w$, the first line of (8.108) therefore gives $|\dot\lambda| \ll_W w^{-1}$. Its second line, interpreted through (8.111) and (8.112), is the exact pair

$$\dot u = u^{(0)} - \dot\lambda q^{\mathrm{reg}}, \qquad \dot a = a^{(0)} - \dot\lambda a_Z.$$

Consequently, prime by prime,

$$\dot u_{j(p)} + \dot\lambda g_p = u_{j(p)}^{(0)} + \dot\lambda(g_p - q_{j(p)}^{\mathrm{reg}}) = u_{j(p)}^{(0)} + \dot\lambda c_p.$$

Equations (8.113), (8.114), and the bounds for $\dot\lambda, a^{(0)}$ now prove

$$|\dot\lambda| \ll_W w^{-1}, \qquad \max_{W<p\le y} |\dot u_{j(p)} + \dot\lambda g_p| + \|\dot a\| + w|\dot\lambda| \le C_*. \tag{8.115}$$

The leading continuum blocks give a constant $C_{\mathrm{inv},W}$ independent of the tilt box and the regular mesh. Before this constant or the ODE box is introduced, W has already made the box-independent terms C_{pow}/W, $C_{\mathrm{pow}}^\sharp/W$, and $C_{\mathrm{rel}}(C_{\mathrm{cmp}})/W$ smaller than the fixed continuum gaps. All dependence on a later fixed tilt box is in $\epsilon_{B,W}(n)$, its explicitly rescaled relative version $\epsilon_{B,W,C_{\mathrm{cmp}}}^{\mathrm{rel}}(n)$, or the analogous fixed-mesh remainders, all of which tend to zero with n. Fix $C_* = 2C_{\mathrm{inv},W}$. After the box below is prechosen, only the threshold for n is increased so that all these B_*, W-dependent vanishing errors use the spare factor two. Thus (8.115) holds with this noncircular C_*.

We now remove the possible bootstrap circularity. Define the effective norm

$$\|\xi\|_{\mathrm{eff}} := \max_{W<p\le y} |u_{j(p)} + \lambda g_p| + \|a\| + w|\lambda|.$$

Choose in advance $B_* = 4C_*$, after W but before the ODE. On the prechosen box $\|\xi\|_{\mathrm{eff}} \le B_*$, the logarithm of the Radon–Nikodym derivative is bounded by

$$\frac{B_*}{L} \sum_{p>W} v_p(m) + O_W(B_*/L) \le \frac{B_*}{\log W} + o_W(1).$$

This gives a fixed dominating density. More sharply, the tilted marked lemmas show that the leading h, K blocks are unchanged and that the box-dependent local errors are $O_{B_*,W}(1/L)$. After increasing the fixed-mesh threshold for n, all covariance gaps and row estimates used in (8.115) remain within the reserved factor two. In particular its right side is at most $2C_{\mathrm{inv},W} = C_* < B_*/2$ throughout this box. The box was chosen before the path, so this conclusion is not being used to define its own domain.

The constants C_* and B_* are uniform over all permitted regular meshes. They may therefore be fixed after W and before a particular mesh is instantiated; after the mesh and the box are fixed, only the threshold for n is increased. In particular, the choice of W never depends on B_*. Fix such a mesh and such an n. In the gauge coordinates the parameter space is finite-dimensional and $\mathcal{M}$ is C^∞. On the closed effective box

$$\mathcal{K}_* := \{\xi : \|\xi\|_{\text{eff}} \le B_*\}$$

the preceding covariance gaps make $D\mathcal{M}(\xi)$ invertible and

$$\|D\mathcal{M}(\xi)^{-1}\tau\|_{\text{eff}} \le C_* \qquad (\xi \in \mathcal{K}_*).$$

Solve the straight-target ODE

$$\xi'(t) = D\mathcal{M}(\xi(t))^{-1}\tau, \qquad \xi(0) = 0. \tag{8.116}$$

Let $[0, T_{\max})$ be its maximal solution interval while the path remains in the interior of $\mathcal{K}_*$. Local existence follows from (8.107). More explicitly, $D\mathcal{M}$ is smooth and invertible on a neighborhood of the closed box, so $\xi \mapsto D\mathcal{M}(\xi)^{-1}\tau$ is locally Lipschitz there; the finite-dimensional Picard–Lindelöf theorem gives both local existence and uniqueness. For $t < \min(T_{\max}, 1)$, integration gives

$$\|\xi(t)\|_{\text{eff}} \le C_* t \le C_* = B_*/4.$$

If $T_{\max} \le 1$, the same derivative bound makes $\xi(t)$ Cauchy as $t \uparrow T_{\max}$. Its limit lies in the strictly smaller box $\|\xi\|_{\text{eff}} \le B_*/4$, where the Jacobian remains invertible. The local existence theorem therefore restarts the solution past $T_{\max}$, a contradiction. Hence $T_{\max} > 1$. The chain rule gives

$$\frac{d}{dt}\mathcal{M}(\xi(t)) = \tau, \qquad \mathcal{M}(\xi(1)) = \tau.$$

To check the converse coordinate implication, let Δ^{out} be the resulting vector of band-moment increments. Equality against every $v \in \mathcal{G}$ implies

$$\Delta^{\text{out}} - \Delta = cD\alpha$$

for some scalar c, because the ordinary annihilator of $\mathcal{G}$ is $\operatorname{span}(D\alpha)$. Preservation of the Z-moments preserves the physical and head moments, so the exact logarithmic identity turns the slow equation into $\alpha^T\Delta^{\text{out}} = \alpha^T\Delta$. Hence $c\,\alpha^T D\alpha = 0$, and $c = 0$. Thus $\mathcal{M}(\xi(1)) = \tau$ is precisely all band, head, physical, and mass equations. Existence, rather than strict convexity, has now been proved; strict convexity only supplies uniqueness in the fixed gauge.

For later use we record the primewise row bound for every block of this inverse. Along (8.116),

$$\begin{aligned} \left|\operatorname{Cov}_\xi(v_p, \dot{\mathscr{Q}}_{\text{quot}})\right| &\ll_W p^{-1}, \\ \left|\operatorname{Cov}_\xi(v_p, \dot{a}^T Z)\right| &\ll_W (pL)^{-1}, \\ \left|\dot{\lambda}\operatorname{Cov}_\xi(v_p, C_g)\right| &\ll_W w^{-1}\frac{w}{p} \ll_W p^{-1}. \end{aligned} \tag{8.117}$$

For the first line one must include the leading row, not only its arithmetic error. If u_j are the bounded effective quotient coefficients, then

$$\left|\operatorname{Cov}_\xi\left(v_p, \sum_j u_j\Omega_j\right)\right| \le \frac{C_W}{p} + C_W\frac{t_p}{p}\sum_{W<q\le y}\frac{t_q}{q} + o(1/p) \ll_W \frac{1}{p}.$$

This uses the diagonal multiplier, the product kernel (8.11), and the box-uniform full-valuation transfer (8.47). The second line follows from the full head/physical marked family (8.100) (whose physical coordinate is (7.61)), and the third from (8.78). Multiplication by $q_n/L \asymp N/L$ and integration over $0 \le t \le 1$ show that the bridge changes every marked moment by $O_W(N/(pL))$. Combined with (6.55), this proves (8.105).

Only z^{act} was reweighted. The effective box bounds its density distortion between two fixed positive constants. Hence, for a constant $C_{B_*,W}$ independent of m, n,

$$0 \le z_m(\xi(1)) \le C_{B_*,W} z_m^0 \le \frac{C'_{B_*,W}}{L}.$$

On a shared clean coordinate the only frozen summand is the protected layer, also $O_W(1/L)$; the top support is disjoint. Therefore the complete coordinate weight is between zero and $(C'_{B_*,W} + C''_W)/L < 1$ for all sufficiently large n. The positive zero-head pool persists. The disjoint top support and the frozen additive protected summand are untouched, proving the remaining margin assertions. $\square$

Corollary 8.8 (Mesh-uniform marked-rate constant)**.** *After W is fixed there exist $\eta_0 > 0$ and a finite constant $C_{\tan}$, chosen before the particular fine mesh and before $w = \delta + \eta$, with the following property. For every permitted regular mesh with $w < \eta_0$, once that mesh is fixed there is a threshold $n_0 = n_0(W, \mathcal{P}, B_*)$ such that, for all $n \ge n_0$, the post-bridge residual satisfies*

$$|r_p| \le C_{\tan} \frac{N}{pL} \qquad (W < p \le y).$$

The threshold is not asserted to be uniform over an infinite family of meshes, but the displayed constant is.

Proof. Choose η_0 below the continuum inverse thresholds in Lemma 8.4 and Lemma 8.5. The quotient and nuisance rows in (8.117) are $O_W(1/p)$, with constants independent of the permitted mesh. The slow row is uniform for the nontrivial reason

$$|\dot{\lambda}|\,|\operatorname{Cov}_\xi(v_p, C_g)| \ll_W \frac{1}{w}\frac{w}{p} \ll_W \frac{1}{p},$$

so its constant is also independent of w. Integrating the three rows for $0 \le t \le 1$ and multiplying by $q_n/L \asymp N/L$ gives a mesh- and w-independent addition to the rough-stage marked-rate constant. Enlarge the previously fixed $C_{\tan}$ by this amount. All box- and mesh-dependent errors tend to zero after the mesh is fixed and are absorbed by increasing only n_0. This proves the stated quantifiers. $\square$

8.7 Choice of constants and the bridge output

The order of constants is important. First fix the tangent comparison ratios $r_0 < 3/2$ and $\rho > 1$ with $\rho^3 < r_0$, together with the anchor and bank data. Next choose one sufficiently large, fixed, nonprime W containing the anchor support, put every prime at most W in the finite head, and make the box-independent quantities

$$\frac{C_{\mathrm{sf}} + C_{\mathrm{pow}} + C^\sharp_{\mathrm{pow}} + C_{\mathrm{rel}}(C_{\mathrm{cmp}})}{W}$$

smaller than the reserved continuum covariance and inverse gaps. Never enlarge W afterward. The compact ODE box is chosen later; its dependence appears only in remainders tending to zero with n. With this head fixed, choose the rough-selector constants and the protected/active split, then choose $\delta_* < 1/18$ and construct the guarded selector. Record its surviving slack σ and fixed-rate constant $C_{\tan}$. Only then choose the fixed regular band parameters δ, η small enough for the bridge and tangent inequalities. Finally fix that mesh and let $n \to \infty$, with its threshold allowed to depend on all preceding fixed choices. This order uses no uniformity for a growing head block.

Combining Sections 8.1 and 8.2 and Proposition 8.7, the fractional point passed to the tangent stage has all of the following exact properties:

(i) every nontrivial complete rough-signature row has its prescribed integer quota, and the flexible smooth row has the integer quota $q_{\mathrm{sm}}^{\mathrm{flex}}(d)$;

(ii) $r_p = 0$ for $p \le W$ and for $p > y$;

(iii) the physical logarithm is exact, so (8.101) holds;

(iv) every band satisfies (8.104), and every medium-prime residual satisfies (8.105); and

(v) on clean smooth endpoints the still-frozen protected layer and the active ceiling provide two-sided slack σ/L; on clean nonsmooth endpoints the two-sided slack is instead the broad-selector margin left by (6.35).

No integer deletion bracket and no unverified endpoint sign is used in this construction.

9 The finite-band tangent absorber

We now cancel the remaining residual at the primes $W < p \le y$. The inputs from Section 8 are

$$\sum_{p\in\mathcal{P}_j} r_p = 0, \qquad |r_p| \le C_{\tan}\frac{N}{pL}, \qquad C_{\tan} := C_{1,W}, \tag{9.1}$$

for every fixed exponent band, together with exact complete rough rows, exact head valuations and exact ordinary logarithm. There is also a fixed $\sigma > 0$ such that every clean broad endpoint in an active, nondedicated row has two-sided slack σ/L. In the smooth row its lower margin comes from the frozen protected summand (and its upper margin from the surviving active ceiling); in a nontrivial rough row both margins are the broad-selector margins left after the $o(1/L)$ row correction. The residual convention is always target minus current. All valuations in (9.1) are full valuations, including prime powers.

9.1 Uniform clean common-multiplier lists

Fix $1 < r_0 < 3/2$. For distinct primes

$$W < v \le u \le y, \qquad u/v \le r_0,$$

put

$$I_{uv} := \left(\frac{n}{v}, \frac{2n - Kh}{u}\right].$$

For $a \in I_{uv}$, both va and ua lie in the broad interval $J = (n, 2n - Kh]$. Let $\mathscr{D}_{\mathrm{row}}$ be the set of fully dedicated nonexceptional rows fixed in (6.42). Under the endpoint-deletion convention used here this set is empty; we nevertheless retain the condition below so that the list definition itself records the required active-row test. Use the exhaustive numerical guard set Γ_{num} defined in (6.39). It consists of the modified external anchors, the fixed exceptional upper factors, every withheld bank donor, and both possible states of every bank path. Later bridge and tangent coordinates are flexible and hence are not an additional class of guards. Define

$$\mathcal{L}_{uv}^{+} := \left\{a \in I_{uv} \cap \mathbb{Z} : (a, P_{\mathrm{hd}}) = 1, \quad X_{R_y(a)} \ge X_0, \quad R_y(a) \notin \mathscr{D}_{\mathrm{row}}, \quad ua, va \notin \Gamma_{\mathrm{num}}\right\}. \tag{9.2}$$

Since $u, v \le y$,

$$R_y(ua) = R_y(a) = R_y(va). \tag{9.3}$$

Thus a switch between ua and va stays in one complete rough signature row.

Lemma 9.1 (Uniform common list)**.** *With the exceptional cutoff fixed once by* (6.1)*, there is a fixed* $\kappa > 0$ *such that*

$$|\mathcal{L}_{uv}^{+}| \geq \kappa \frac{n}{u}$$

uniformly for every permitted pair u, v *and every sufficiently large* n*.*

Proof. Put $\delta_{\mathrm{hd}} = \varphi(P_{\mathrm{hd}})/P_{\mathrm{hd}}$. The interval length is

$$|I_{uv}| = \frac{n}{u}\left(2 - \frac{u}{v} - \frac{Kh}{n}\right) + O(1).$$

Counting the fixed reduced residue classes modulo P_{hd} gives

$$\#\{a \in I_{uv} : (a, P_{\mathrm{hd}}) = 1\} = \left\{\delta_{\mathrm{hd}}\left(2 - \frac{u}{v}\right) + o(1)\right\}\frac{n}{u}. \tag{9.4}$$

The error is uniform because $n/u \geq n/y = n^{7/9}$.

We next remove the exceptional rows. Write $a = Rb$, where $R = R_y(a)$ and $P^{+}(b) \leq y$. If $X_R = 2n/R < X_0 = n^{\delta_*}$, then

$$R > \frac{2n}{X_0}, \qquad b < \frac{2X_0}{u}. \tag{9.5}$$

For an upper bound, discard the additional lower cutoff $R > 2n/X_0$. For fixed b, the full physical interval forced by $a \in I_{uv}$ has location and length both $\asymp_{r_0} n/(ub)$; whenever it can contain an exceptional multiplier, (9.5) makes this length $\gg_{r_0} n/X_0$. The interval Selberg sieve (2.11) applies uniformly: its lcm remainder level is $y^4 = n^{8/9}$, while

$$n/X_0 = n^{1-\delta_*}, \qquad \delta_* < 1/18.$$

For a fixed sieve constant $C_{\mathrm{sv}} > 0$, it follows that the number E_{uv} of exceptional multipliers is

$$\begin{aligned} E_{uv} &\ll \frac{n}{u \log y} \sum_{b \leq 2X_0/u} \frac{1}{b} + \frac{X_0}{u} \frac{y^4}{(\log y)^2} \\ &\leq \left\{C_{\mathrm{sv}} \frac{\delta_*}{\theta} + o(1)\right\} \frac{n}{u}. \end{aligned} \tag{9.6}$$

Indeed the second term is $o(n/u)$, since

$$\frac{X_0 y^4}{n(\log y)^2} = n^{\delta_* + 8/9 - 1 + o(1)} = o(1).$$

If $u > 2X_0$, this sum is empty. Notice that the loss is a small fixed fraction, not $o(n/u)$.

It remains to check the actual guards. We give the complete census of those guards that can meet a lower endpoint ua or va. Every modified prefix anchor is Pq, where $P > y$ is its unique rough marker and every prime factor of q is at most $2R + 1 \leq W$; hence it contains no medium label. A row-zero anchor is a prime greater than n, so it contains no medium label either. A promoted anchor is $2^{k_p} p^{e_p}$. If $p > y$, its only smooth prime is 2; if $p \leq y$, it is one of at most $\pi(y)$ promoted smooth anchors. Consequently only $O(\pi(y))$ anchors can be divisible by a prime in $(W, y]$.

There are $O(yL)$ component factors in the fully reserved bank by (5.14); taking both states changes only the implied constant. These are the only remaining lower guards. Indeed, $G_{\mathrm{fix}} \subset (2n, M]$ by (6.37), and every member of E_{donor} lies above $2n$, so neither class can equal a lower endpoint. Thus the number of endpoint-relevant guards is

$$O(\pi(y)) + O(yL) = O(yL).$$

For a fixed label u, the equation $ua = g$ determines at most the one integer multiplier $a = g/u$; the same statement holds for v. Therefore the guard deletion is, uniformly in the permitted pair,

$$O(yL) = o(n/u), \tag{9.7}$$

because $u \le y$ and $y^2L/n = o(1)$. The dedicated-row condition costs zero because $\mathscr{D}_{\mathrm{row}} = \varnothing$. More generally, the quantitative data in (6.42) would bound its loss by $O\{yL(y/u+1)\} = o(n/u)$. Thus the definition cannot call on a row whose flexible quota was set to zero. Terminal exceptional rows were already removed in (9.6). Forced exceptional upper factors lie above $2n$ and cannot occur in a lower list.

Combining (9.4), (9.6), and (9.7), the one-time inequality (6.1) gives

$$C_{\mathrm{sv}}\frac{\delta_*}{\theta} < \frac{1}{4}\delta_{\mathrm{hd}}(2-r_0), \qquad \delta_{\mathrm{hd}}(2-r_0) - C_{\mathrm{sv}}\frac{\delta_*}{\theta} > \frac{3}{4}\delta_{\mathrm{hd}}(2-r_0) > 0.$$

Thus, without changing δ_*, one may take

$$\kappa = \frac{1}{2}\left\{\delta_{\mathrm{hd}}(2-r_0) - C_{\mathrm{sv}}\frac{\delta_*}{\theta}\right\} > 0.$$

□

We record explicitly why every endpoint now retained in (9.2) has usable slack. If $R_y(a) = 1$, the frozen protected smooth summand supplies the lower floor and the fitted active summand stays below its broad ceiling. If $R_y(a) > 1$, the row is nondedicated by definition of the list, and the broad raw floor and ceiling survive the row correction because that correction is $o(1/L)$ per clean coordinate. After decreasing the fixed constant σ, if needed, both cases give

$$\frac{\sigma}{L} \le x_{ua}, x_{va} \le 1 - \frac{\sigma}{L}. \tag{9.8}$$

Thus one pair u, v has at least

$$\kappa\sigma\frac{N}{u} \tag{9.9}$$

real two-sided capacity. This floor is used for tangents only now; it was not a source of bridge covariance.

9.2 Finite-band earthmover estimates

Use the exponent bands corresponding to (8.20):

$$\mathcal{B}_0 = (W, n^{\delta_b}], \qquad \mathcal{B}_j = (n^{a_{j-1}}, n^{a_j}] \quad (1 \le j \le s),$$

where $a_0 = \delta_b$, $a_s = \theta = 2/9$, and

$$c_{\mathrm{mesh}}\eta a_{j-1} \le a_j - a_{j-1} \le \eta a_{j-1}.$$

Set

$$w_{\tan} := \delta_b + \theta\eta.$$

These are the same bands as in the bridge, merely written in the $\log n$-exponent convention.

Choose a fixed $\rho > 1$ with $\rho^3 < r_0$. Partition each band, separately, into multiplicative ρ-cells, merging a possible terminal fragment with its neighbor. Any labels paired within one cell or in adjacent cells then have ratio at most $\rho^3 < r_0$. The prime number theorem, after W is chosen sufficiently large depending on ρ, gives

$$\#\{p \text{ in a cell at scale } Q\} \asymp_\rho \frac{Q}{\log Q}. \tag{9.10}$$

This is a fixed-ratio interval, so no short-interval prime theorem is being used.

Lemma 9.2 (Finite-band transport)**.** *In each exponent band $[a, b]$, there is a nonnegative directed flow (f_{uv}) on its prime labels such that*

$$\sum_v (f_{pv} - f_{vp}) = r_p, \tag{9.11}$$

$$f_{uv} > 0 \implies \max(u, v)/\min(u, v) \le \rho^3 < r_0. \tag{9.12}$$

Writing $f_p := \sum_v (f_{pv} + f_{vp})$, the union of these flows over all bands may be chosen to satisfy

$$\#\{(u, v) : f_{uv} > 0\} = O_\rho(\pi(y)), \tag{9.13}$$

$$\sum_{u,v} f_{uv} \ll_\rho C_{\tan} N w_{\tan} + o(N), \tag{9.14}$$

$$\frac{p f_p}{N} \ll_\rho C_{\tan} \left\{ t \log \frac{b}{t} + \frac{1}{L} \right\} \quad (p \asymp_\rho n^t), \tag{9.15}$$

$$\sup_p \frac{p f_p}{N} \ll_\rho C_{\tan} w_{\tan} + o(1). \tag{9.16}$$

All constants are uniform over the finitely many bands in the permitted mesh.

Proof. Fix one band and write its consecutive prime cells as $\mathcal{Q}_1, \dots, \mathcal{Q}_m$. Put

$$s_i := \sum_{p \in \mathcal{Q}_i} r_p, \qquad F_i := \sum_{h \le i} s_h \quad (1 \le i < m), \qquad F_0 = F_m = 0.$$

The last equality is exactly the band-balance assertion in (9.1). If the boundary between $\mathcal{Q}_i$ and $\mathcal{Q}_{i+1}$ is $Q = n^t$, band balance also gives

$$|F_i| = \left| \sum_{Q < p \le n^b} r_p \right|.$$

Using the pointwise bound in (9.1) and subtracting the two instances of (2.9), at their actual endpoints, we obtain uniformly

$$|F_i| \ll C_{\tan} \frac{N}{L} \left\{ \log \frac{b}{t} + \frac{1}{tL} \right\}. \tag{9.17}$$

The term $1/(tL)$ is the endpoint error $O(1/\log Q)$; this remains valid in the moving first band, where $t \ge \log W/L$. Thus no fixed-positive lower exponent is being assumed.

We next give the promised flow algorithm. Across boundary i, send the signed load F_i from $\mathcal{Q}_i$ to $\mathcal{Q}_{i+1}$: if $F_i < 0$, reverse the direction. On either side distribute its absolute value uniformly among all primes of that cell. The two resulting lists of nonnegative port masses have the same total $|F_i|$; the usual greedy matching of their first still-positive entries realizes this transport with at most $|\mathcal{Q}_i| + |\mathcal{Q}_{i+1}| - 1$ positive edges.

Let d_p be the boundary-flow divergence at p, with outgoing minus incoming convention. Since the signed boundary load on the right of cell i is F_i, while that on its left is F_{i-1},

$$\sum_{p \in \mathcal{Q}_i} d_p = F_i - F_{i-1} = s_i.$$

Consequently $q_p := r_p - d_p$ has sum zero in each cell. Greedily match the positive and negative q_p's inside that cell, directing flow from the former to the latter. This uses at most $|\mathcal{Q}_i| - 1$ positive edges and gives divergence exactly q_p. Adding the boundary and internal flows proves (9.11). Endpoints of an internal edge lie in one cell, and endpoints of a boundary edge lie in adjacent cells. The possible merged terminal cell was allowed for in the exponent ρ^3, so (9.12) follows.

This construction also proves the quantitative claims without an unproved distribution step. Put $m_i := |\mathcal{Q}_i|$. The total boundary traffic incident to $p \in \mathcal{Q}_i$ is at most

$$b_p := \frac{|F_{i-1}| + |F_i|}{m_i}.$$

The internal traffic incident to p is $|q_p|$, and $|d_p| \le b_p$; hence

$$f_p \le |r_p| + 2b_p. \tag{9.18}$$

For $p \asymp_\rho Q = n^t$, (9.10) says $m_i \asymp_\rho p/(tL)$. Substitution of (9.17) into (9.18), together with $|r_p| \le C_{\tan} N/(pL)$, proves (9.15).

There are only a bounded number of greedy matchings incident to a cell: one internal matching and at most two boundary matchings. Summing their edge counts proves (9.13). Moreover, the cross-boundary traffic is $\sum_i |F_i|$, while the internal traffic is

$$\frac{1}{2}\sum_p |q_p| \le \frac{1}{2}\sum_p |r_p| + \sum_i |F_i|.$$

Therefore

$$\sum_{u,v} f_{uv} \le \frac{1}{2}\sum_{W<p\le y} |r_p| + 2\sum_{\text{boundaries}} |F_i|. \tag{9.19}$$

The first term is

$$\frac{1}{2}\sum_{W<p\le y} |r_p| \ll C_{\tan}\frac{N\log L}{L} = o(N)$$

by (2.9). Consecutive cell boundaries have exponent spacing $\asymp_\rho 1/L$, so (9.17) and an upper Riemann sum give

$$\begin{aligned}\sum_{\text{boundaries}} |F_i| &\ll_\rho C_{\tan} N\left\{\int_0^{\delta_b}\log\frac{\delta_b}{t}\,dt + \sum_{j=1}^{s}\int_{a_{j-1}}^{a_j}\log\frac{a_j}{t}\,dt\right\} + o(N)\\ &\ll_\rho C_{\tan} N(\delta_b + \theta\eta) + o(N).\end{aligned} \tag{9.20}$$

Here the accumulated $1/(tL)$ endpoint terms are $O(N\log L/L) = o(N)$. Also $\int_0^{\delta_b}\log(\delta_b/t)\,dt = \delta_b$, while

$$\int_a^b \log(b/t)\,dt \le \frac{(b-a)^2}{a} \le \eta(b-a)$$

for every regular positive band; summing gives $O(\theta\eta)$. Equations (9.19)–(9.20) prove (9.14). Finally, $t\log(b/t) \le b - t$; it is at most δ_b in the first band and at most $\theta\eta$ in a positive regular band. This proves (9.16) and completes the construction. □

At a boundary $Q = n^t$, (9.10) and (9.9) give aggregate port capacity

$$\asymp_\rho \kappa\sigma\frac{N}{\log Q} = \asymp_\rho \kappa\sigma\frac{N}{tL}.$$

Comparing this with (9.17) shows explicitly that the maximum boundary utilization is

$$\ll_\rho \begin{cases} C_{\tan}\delta_b/(\kappa\sigma), & \mathcal{B}_0,\\ C_{\tan}\theta\eta/(\kappa\sigma), & \mathcal{B}_j,\ j\ge 1.\end{cases} \tag{9.21}$$

9.3 Requests and simultaneous collision avoidance

For a directed support edge $e = (s(e), t(e))$, put

$$f_e := f_{s(e)t(e)}, \qquad U_e := \max\{s(e), t(e)\}, \qquad V_e := \min\{s(e), t(e)\}.$$

The common list is indexed by the unordered pair (U_e, V_e), while the source–sink orientation is retained as part of the request. For $f > 0$, put

$$k(f) := \max\left\{1, \left\lceil \frac{4Lf}{\sigma} \right\rceil\right\}. \tag{9.22}$$

Split an edge e with $f_e > 0$ into $k(f_e)$ equal requests. The definition gives

$$k(f) \geq \frac{4Lf}{\sigma}, \qquad \frac{f}{k(f)} \leq \frac{\sigma}{4L}, \qquad f \leq \frac{\sigma k(f)}{4L}.$$

Thus every request has size at most $\sigma/(4L)$, and

$$k(f) \leq \frac{4Lf}{\sigma} + 1. \tag{9.23}$$

The second inequality is essential: a small-piece assertion alone would not control the number of independent random choices.

Let k_p be the number of requests incident to the prime label p, and let K_{req} be the total number of requests. By (9.13),

$$\begin{aligned} k_p &\leq \frac{4L}{\sigma} f_p + \deg_{\mathrm{supp}}(p), \\ K_{\mathrm{req}} &\leq \frac{4L}{\sigma} \sum_{u,v} f_{uv} + O(\pi(y)). \end{aligned}$$

Equations (9.14), (9.15), and (9.16) imply

$$\begin{aligned} \frac{pk_p}{n} &\ll_\rho \frac{C_{\tan}}{\sigma} w_{\tan} + o(1), \\ \frac{K_{\mathrm{req}}}{n} &\ll_\rho \frac{C_{\tan}}{\sigma} w_{\tan} + o(1). \end{aligned} \tag{9.24}$$

Indeed, the direct term contributes $O(C_{\tan} \log L/(\sigma L)) = o(1)$. The $+1$ terms in (9.23) are harmless because

$$\frac{\pi(y)}{n} = o(1), \qquad \frac{y\pi(y)}{n} = n^{2\theta-1+o(1)} = o(1).$$

In particular, if

$$\frac{C_{\tan}}{\kappa\sigma} w_{\tan} \tag{9.25}$$

is sufficiently small, every prime port retains a fixed capacity margin. Under the stronger smallness condition imposed in (9.33), the implicit constant in (9.24) may also be absorbed so that, uniformly for every directed support edge e,

$$k(f_e) \leq k_{U_e} \leq \frac{\kappa}{4} \frac{n}{U_e} \leq \frac{1}{4} |\mathcal{L}^+_{U_e V_e}|.$$

Moreover, this request bound implies the real edge-capacity bound

$$f_e \leq \frac{\sigma k(f_e)}{4L} \leq \frac{\kappa\sigma}{16} \frac{n}{U_e L} = \frac{\kappa\sigma}{16} \frac{N}{U_e}.$$

On the other hand, the actual two-sided capacity of the common list is

$$|\mathcal{L}^{+}_{U_eV_e}|\frac{\sigma}{L} \geq \kappa\frac{n}{U_e}\frac{\sigma}{L} = \kappa\sigma\frac{N}{U_e}.$$

Hence each real edge uses at most one sixteenth of its literal available capacity. Even before the random experiment, it also has fewer requests than list entries. The later local lemma enforces the stronger global endpoint disjointness.

Let $\mathscr{R}$ be the finite request set. For a request $\mathfrak{r}$ on labels $\{U_\mathfrak{r}, V_\mathfrak{r}\}$, with $U_\mathfrak{r} \geq V_\mathfrak{r}$, let $Z_\mathfrak{r}$ be uniform on $\mathcal{L}^{+}_{U_\mathfrak{r}V_\mathfrak{r}}$. Its directed source and sink are also retained from the flow edge that created it. Choose all the variables $Z_\mathfrak{r}$ independently. Its two numerical endpoints are

$$E_\mathfrak{r} := \{U_\mathfrak{r}Z_\mathfrak{r}, V_\mathfrak{r}Z_\mathfrak{r}\}.$$

They are distinct within one request because $U_\mathfrak{r} \neq V_\mathfrak{r}$. For every unordered pair of distinct requests define the bad event

$$\mathscr{E}_{\mathfrak{r},\mathfrak{s}} := \{E_\mathfrak{r} \cap E_\mathfrak{s} \neq \varnothing\}. \tag{9.26}$$

This definition includes all four cross-request endpoint equalities.

Suppose first that the two label pairs are disjoint. For fixed endpoint labels α and β, the equation $\alpha a = \beta b$ has, because the two primes are distinct,

$$a = \beta m, \qquad b = \alpha m.$$

The endpoint range gives $m = O(n/(\alpha\beta))$; the harmless literal $+1$ is absorbed because $\alpha\beta \leq y^2 = o(n)$. Moreover the larger label in the first request is at most $r_0\alpha$, and similarly for the second. Lemma 9.1 therefore gives a product of list sizes at least $\kappa^2 n^2/(r_0^2\alpha\beta)$. Taking the union over the four endpoint equations yields

$$\mathbb{P}(\mathscr{E}_{\mathfrak{r},\mathfrak{s}}) \ll_{r_0} \frac{1}{\kappa^2 n}. \tag{9.27}$$

Now suppose that the requests share a label p. After writing their pairs as $\{p, r\}$ and $\{p, q\}$, the four equalities are

$$pa = pb, \qquad pa = qb, \qquad ra = pb, \qquad ra = qb.$$

For $pa = pb$, there are at most $O_{r_0}(n/p)$ possible common multipliers, while both lists have size $\gg_{r_0} \kappa n/p$. Its probability is therefore $O_{r_0}(p/(\kappa^2 n))$. Each equality with distinct endpoint labels has $O(n/(\alpha\beta))$ raw solutions as above and probability $O_{r_0}(1/(\kappa^2 n))$. This also covers identical label pairs: there are then two same-label equalities, with the same bound. Hence

$$\mathbb{P}(\mathscr{E}_{\mathfrak{r},\mathfrak{s}}) \ll_{r_0} \frac{p}{\kappa^2 n}. \tag{9.28}$$

These cases exhaust the event in (9.26).

Fix a request with labels u, v. There are at most k_u other requests sharing u, at most k_v sharing v, and at most K_{req} remaining requests. Thus

$$\sum_{\mathfrak{s}\neq\mathfrak{r}} \mathbb{P}(\mathscr{E}_{\mathfrak{r},\mathfrak{s}}) \ll_{r_0} \frac{1}{\kappa^2}\left(\frac{uk_u + vk_v}{n} + \frac{K_{\mathrm{req}}}{n}\right).$$

Equations (9.24), (9.27), and (9.28) therefore give

$$B_{\mathrm{coll}} := \sup_{\mathfrak{r}} \sum_{\mathfrak{s}\neq\mathfrak{r}} \mathbb{P}(\mathscr{E}_{\mathfrak{r},\mathfrak{s}}) \ll_{r_0,\rho} \frac{C_{\tan}}{\kappa^2\sigma} w_{\tan} + o(1). \tag{9.29}$$

Thus every bad event, including every repeated-label case, is present in the displayed probability ledger. Choose the fixed band parameters so that $B_{\mathrm{coll}} \leq 1/8$.

For completeness, we now specify the dependency graph and verify the factor-four bookkeeping. Join two bad events when their unordered request pairs intersect. An event is a function of its two request variables, and the underlying variables are mutually independent; therefore it is independent of the sigma-algebra generated by all nonneighbors. For an event E, its neighborhood lies in the union of the bad events involving either of its two requests, and hence

$$\sum_{E' \sim E} \mathbb{P}(E') \leq 2B_{\mathrm{coll}}.$$

Uniformly, $\mathbb{P}(E) = O_{r_0}(y/(\kappa^2 n)) = o(1)$, so $x_E := 2\mathbb{P}(E)$ belongs to $[0, 1)$ for sufficiently large n. Also

$$\sum_{E' \sim E} x_{E'} \leq 4B_{\mathrm{coll}} \leq \frac{1}{2}.$$

Using $\prod_i (1 - z_i) \geq 1 - \sum_i z_i$ for nonnegative z_i with sum at most one,

$$x_E \prod_{E' \sim E} (1 - x_{E'}) \geq 2\mathbb{P}(E) \left(1 - \sum_{E' \sim E} x_{E'}\right) \geq \mathbb{P}(E).$$

The asymmetric local lemma from Section 2.5 therefore gives a simultaneous choice with no endpoint collision. Since every list was guarded in advance, these endpoints are also distinct from every anchor, fixed factor, bank state, and withheld donor.

9.4 Exact flow algebra and preserved invariants

Lemma 9.2 constructs the directed flow, and the request splitting preserves its divergence exactly. If $s(e)$ and $t(e)$ denote respectively the source and sink label of a request and $\tau_e > 0$ its size, then

$$\sum_e \tau_e (\mathbf{e}_{s(e)} - \mathbf{e}_{t(e)}) = r. \tag{9.30}$$

This equation fixes the sign convention: the selector change is $+r$. Equivalently, graph boundary is source minus sink and equals r.

For the multiplier chosen for a request, use the signed coordinate move

$$\tau_e (\mathbf{1}_{s(e)a(e)} - \mathbf{1}_{t(e)a(e)}).$$

Its full valuation change is

$$v(s(e)a(e)) - v(t(e)a(e)) = \mathbf{e}_{s(e)} - \mathbf{e}_{t(e)}. \tag{9.31}$$

This remains true even if the multiplier contains $s(e)$, $t(e)$, or higher powers of either, because the complete vector $v(a(e))$ cancels.

The moves have the following exact invariants.

(i) By (9.3), each move stays inside one complete rough-signature row; cardinality and every valuation above y are unchanged.

(ii) The multiplier is P_{hd}-free and $s(e), t(e) > W$, so all head valuations are unchanged.

(iii) Both labels lie in one exponent band, so every exact band total is unchanged.

(iv) Equation (9.30) and (9.31) make every medium-prime valuation exact.

The preliminary selector was exact in ordinary logarithm. Since its head and rough coordinates were already exact, the residual convention gives

$$\sum_{W<p\le y} r_p \log p = 0.$$

Indeed, the target and the current selector have the same valuations outside $(W, y]$, so the logarithm of their product ratio is exactly the left side. Taking the prime-log scalar product of (9.30) now gives the complete calculation

$$\Delta \log = \sum_e \tau_e\{\log s(e) - \log t(e)\} = \sum_{W<p\le y} r_p \log p = 0.$$

Thus no separate logarithmic balance is required inside an exponent band.

All chosen numerical endpoints are distinct, and every request has size at most $\sigma/(4L)$. Thus each coordinate is changed at most once and (9.8) gives the explicit post-move margin

$$\frac{3\sigma}{4L} \le x_a^{\text{new}} \le 1 - \frac{3\sigma}{4L} \qquad \text{for every endpoint used by the tangent.}$$

Every unused coordinate keeps its previous feasible value, so the whole vector remains in $[0, 1]$. At this point the protected layer has been released: the object passed to rounding is the single total vector (x_a), not a separate floor plus an independently roundable selector.

Proposition 9.3 (Exact post-tangent certificate)**.** *Let $\mathcal{A}_R$ denote the remaining flexible candidates of actual complete signature $R = R_y(a)$; thus $R = 1$, not 0, is the smooth row. After the tangent construction,*

$$\boxed{\begin{array}{c} 0 \le x_a \le 1 \quad (a \in \mathcal{A}), \\ \displaystyle\sum_{a\in\mathcal{A}_R} x_a = q_R \in \mathbb{Z} \quad (R \ne 1), \qquad \sum_{a\in\mathcal{A}_1} x_a = q_{\text{sm}}^{\text{flex}}(d) \in \mathbb{Z}, \\ \displaystyle v(B_0) + \sum_{a\in\mathcal{A}} x_a v(a) = v(Y) \qquad \textit{at every prime.} \end{array}} \tag{9.32}$$

Proof. The row identities are unchanged by the tangent and were integer identities before it. Primes $p \le W$ were exact after the head fit, primes $W < p \le y$ are exact by (9.30), and primes $p > y$ were exact by the complete rough-signature equations. This proves the valuation identity simultaneously at all primes. Coordinate feasibility was proved above. $\square$

Since every valuation on the left of (9.32) is nonnegative, the same identity also re-verifies $B_0 \mid Y$ for the *combined* charged base. It is not obtained by multiplying separate divisibility claims. The certificate (9.32) is the precise input to the deterministic rounding and bank absorption in the next section.

9.5 Final parameter order

We summarize the noncircular choice. Fix $r_0 < 3/2$, then fix $\rho > 1$ with $\rho^3 < r_0$. Choose the single finite head cutoff W large enough for the head cells, marked estimates, uniform bridge perturbation, guard estimates, and the fixed-ratio cell prime counts. For the bridge perturbation this choice uses only the box-independent leading constants in (7.37), (7.60), and (8.69); it does not use the radius of the later ODE box. Use the one-time cutoff δ_* already fixed in (6.1), construct the guarded rough selector and the head/physical baseline, and apply Lemma 9.1; this fixes κ and the surviving endpoint margin σ. The uniform estimates of Proposition 8.7 now supply a mesh threshold and a constant $C_{\tan}$ valid for every sufficiently fine regular partition; at this point no particular partition has yet been instantiated. Choose a fixed

$\varepsilon_{\tan} = \varepsilon_{\tan}(r_0, \rho, W, \kappa, \sigma) > 0$ small enough to absorb the fixed constants in (9.21), (9.24), and (9.29). Then choose fixed $\delta_b, \eta > 0$ inside the uniform mesh range and so small that

$$C_{\tan}(\delta_b + \theta\eta) \le \varepsilon_{\tan}\kappa^2\sigma. \tag{9.33}$$

Instantiate the bridge on this partition and then run the tangent. The displayed condition implies both the port condition (9.25) and the local-lemma condition (9.29); request sizes themselves are bounded by the previously fixed slack in (9.22). Finally take n sufficiently large to absorb every accompanying $o(1)$ term, including all B_*, W-dependent power-transfer remainders, and then let $n \to \infty$. Neither W nor any already recorded W-dependent constant is changed afterward.

10 Deterministic rounding and final assembly

We now convert the exact fractional selector into an exact subset of legal factors. We first recall the chronology of the charge, because it rules out a possible circularity. Before the rough selector was solved, the fixed residual factors and the chosen zero states of the universal bank were combined into

$$B_0 = \prod_{a \in G_{\mathrm{fix}}} a \prod_{a \in G^0_{\mathrm{bank}}} a \tag{10.1}$$

as in (6.38). The nonnegative exceptional estimate (6.47), the bank and guard estimate

$$v_p(B_0/B_{\mathrm{exc}}) = o_W\left(\frac{N}{pL}\right) \qquad (p \le y),$$

and the uniform tail supply

$$v_p(T_n) = \frac{(c + o(1))N}{p-1}$$

were compared coordinatewise. At primes in F_{anc}, this comparison used the strict reserve remaining after D was changed to D'; away from F_{anc}, the divisor D' has no valuation. After δ_* was fixed sufficiently small, these inequalities held simultaneously at every $p \le y$. Above y, every fixed or bank-base token was injected into a different withheld tail token with the same complete signature. Thus, before any fractional solve,

$$D'B_0 \mid T_n, \qquad B_0 \mid Y := T_n/D'.$$

This is the combined certificate (6.50); it is not obtained by multiplying separate divisors of T_n.

Let $\mathcal{A} = \coprod_S \mathcal{A}_S$ be the flexible lower candidates remaining after all declared guards. The tangent stage supplies one total weight $x_a \in [0,1]$ for each $a \in \mathcal{A}$ and the exact post-tangent certificate

$$\sum_{a \in \mathcal{A}_S} x_a = q_S \in \mathbb{Z} \qquad \text{for every flexible complete-signature row } S, \tag{10.2}$$

$$v(B_0) + \sum_{a \in \mathcal{A}} x_a v(a) = v(Y) \qquad \text{at every prime.} \tag{10.3}$$

These are precisely the equations in (9.32). In the trivial signature row, the integer in (10.2) is $q^{\mathrm{flex}}_{\mathrm{sm}}(d)$. It is the total flexible smooth-row quota from (8.3). The bridge-active quantity $q^{\mathrm{act}}(d)$ is a real mass and is not used as a rounding equation. Thus the distinction between the integer total row and the real active normalization survives intact into the present argument.

10.1 Column-sparse floating rounding

The complete column family used below is

$$\mathcal{C} := \{C_{p,j} : p \text{ prime},\ j \geq 1,\ p^j \leq M\}, \qquad C_{p,j}(a) := \mathbf{1}_{\{p^j \mid a\}}. \tag{10.4}$$

It is finite: necessarily $p \leq M$ and $1 \leq j \leq \lfloor \log_2 M \rfloor$. Because $M < 3n$ for large n, a legal factor belongs to at most

$$\sum_p v_p(a) = \Omega(a) \leq \log_2 a \leq d_n$$

of these columns.

Lemma 10.1 (Floating rounding)**.** *Suppose a finite set $\mathcal{A}$ is partitioned into rows $\mathcal{A}_S$, $0 \leq x_a \leq 1$, and $\sum_{a \in \mathcal{A}_S} x_a$ is an integer in every row. If every coordinate belongs to at most d members of a finite family of zero-one columns, there are $X_a \in \{0, 1\}$ such that*

$$\sum_{a \in \mathcal{A}_S} X_a = \sum_{a \in \mathcal{A}_S} x_a \qquad \textit{for every } S, \tag{10.5}$$

$$\left| \sum_{a \in \mathcal{A}} (X_a - x_a) C(a) \right| \leq 4d \qquad \textit{for every column } C. \tag{10.6}$$

Proof. Freeze every coordinate as soon as it becomes integral. At an intermediate stage let F be the set of strictly fractional coordinates and put $m = |F|$. A row meeting F contains at least two members of F: after the frozen integral coordinates in that row are removed, the sum of the remaining coordinates is an integer, which one strictly fractional coordinate cannot supply. Hence at most $m/2$ row equations are active.

Retain every active row equation and every column having more than $4d$ members of F. There are at most dm coordinate–column incidences, so fewer than $m/4$ column equations are retained. The retained system is a finite real linear system and has rank less than $3m/4 < m$. It therefore has a nonzero null vector $z = (z_a)_{a \in F}$. Its complete feasible movement interval is

$$I_z := \{t \in \mathbb{R} : 0 \leq x_a + t z_a \leq 1 \text{ for every } a \in F\}.$$

Because every x_a, $a \in F$, is strictly between zero and one, 0 is an interior point of I_z. Since $z \neq 0$, this is a bounded closed interval with two finite endpoints. Move to either endpoint. Every retained equation is preserved and at least one additional coordinate reaches 0 or 1; it is then frozen. Consequently the procedure stops after at most $|\mathcal{A}|$ iterations with an integral vector.

Row equations are never dropped, which proves (10.5). At the first stage when a column is not retained, it has at most $4d$ fractional coordinates. Its value was preserved at every earlier stage, and only those then-fractional coordinates can subsequently change. Each changes by at most one, so its total change from the original vector is at most $4d$. A column retained through the last stage is preserved exactly. This proves (10.6). □

Apply the lemma to the complete-signature rows, with the prime-power columns (10.4) and $d = d_n$. Define the rounded-minus-fractional valuation error

$$e_p := \sum_{a \in \mathcal{A}} (X_a - x_a) v_p(a). \tag{10.7}$$

It is an integer. Indeed, by (10.3), $\sum_a x_a v_p(a) = v_p(Y) - v_p(B_0)$ is an integer, while the rounded sum is also integral. Since

$$v_p(a) = \sum_{j \geq 1} C_{p,j}(a),$$

(10.6) gives

$$|e_p| \le 4d_n \left\lceil \frac{\log(3n)}{\log p} \right\rceil = \beta_p. \tag{10.8}$$

The complete-row equations do essential additional work. If $P > y$, then $v_P(a) = S_P$ is constant on $\mathcal{A}_S$; hence

$$e_P = \sum_S S_P \sum_{a \in \mathcal{A}_S} (X_a - x_a) = 0 \tag{10.9}$$

by (10.5). Thus e is supported on $p \le y$ and lies in the full universal box reserved in Section 5.

10.2 The guarded exactification lemma

For clarity, we isolate the exact hypotheses used in the last discrete step.

Proposition 10.2 (Guarded integral exactification)**.** *Assume the following.*

1. *The flexible candidates are partitioned by their complete signatures, all their row sums are the integers in* (10.2)*, and their fractional weights lie in* $[0, 1]$*.*
2. *The fixed factors, flexible candidates, both states of every bank path, donors, and external anchors satisfy the joint guarding and distinctness conditions of Section 5.*
3. *The charged base is exactly* (10.1)*, and the valuation certificate* (10.3) *holds after that charge.*
4. *Every path satisfies the all-factor row identity* (5.2)*; every base-state row token replaced a distinct withheld donor token before the row quotas were frozen; and the resulting quotas are nonnegative.*
5. *The bank is the two-sided β-universal reserve of Lemma 5.3.*

Then there are $X_a \in \{0, 1\}$ and state indices $\varepsilon_g \in \{0, 1\}$, $g \in \mathcal{B}$, such that, on putting

$$G_{\text{bank}}^{\text{fin}} := \bigsqcup_{g \in \mathcal{B}} G_g^{\varepsilon_g},$$

the three sets G_{fix}, $G_{\text{bank}}^{\text{fin}}$, and $\{a \in \mathcal{A} : X_a = 1\}$ are pairwise disjoint, the flexible row quotas are preserved, and the product of their union is exactly Y.

Proof. Use Lemma 10.1. Equations (10.8)–(10.9) put the integer error e in the universal box. Apply Lemma 5.3 to $-e$. It supplies a subcollection of whole paths to toggle. Set $\varepsilon_g = 1$ exactly for those paths and $\varepsilon_g = 0$ otherwise. The operation is the state replacement

$$G_{\text{bank}}^{0} = \bigsqcup_{g \in \mathcal{B}} G_g^{0} \quad \rightsquigarrow \quad G_{\text{bank}}^{\text{fin}} = \bigsqcup_{g \in \mathcal{B}} G_g^{\varepsilon_g};$$

no bank factor is multiplied onto the already charged base. An untoggled path remains literally in its base state, whereas a toggled path is replaced once, simultaneously in all of its components.

By universality and the definition of a path change,

$$v(G_{\text{bank}}^{\text{fin}}) - v(G_{\text{bank}}^{0}) = \sum_{\varepsilon_g = 1} \{v(G_g^1) - v(G_g^0)\} = -e. \tag{10.10}$$

Likewise, for every complete signature S,

$$m_S(G_{\text{bank}}^{\text{fin}}) - m_S(G_{\text{bank}}^{0}) = \sum_{\varepsilon_g = 1} \{m_S(G_g^1) - m_S(G_g^0)\} = 0$$

by (5.2). Thus every fixed row token charged in the base remains one fixed row token after exactification, and the flexible row equations preserved by floating rounding remain valid.

The rounded flexible change is $+e$, while (10.10) is $-e$. Starting with (10.3), their sum leaves every prime valuation equal to $v(Y)$. Finally, both possible states of every path were jointly guarded before the choice of the toggle subplan. Therefore an arbitrary subplan, including the one just selected, is collision-free internally and disjoint from G_{fix}, the anchors, and all flexible candidates. This proves both the product identity and the asserted distinctness. $\square$

No coordinatewise relative-interior hypothesis is used in this proof. Such interiority was useful upstream in constructing the fractional point, but floating rounding and the universal bank require only $0 \le x_a \le 1$, integer row sums, and the exact charged certificate.

10.3 The final residual set

Theorem 10.3 (Upper-bound construction)**.** *Fix $c > C_0$ and put*

$$M = 2n + \left\lceil c \frac{n}{\log n} \right\rceil .$$

For every sufficiently large n, there is a set of distinct integers

$$\mathcal{S} \subset (n, M]$$

such that

$$\prod_{a \in \mathcal{S}} a = Q(n, M).$$

Proof. We first discharge, in order, the five hypotheses of Proposition 10.2.

(i) Proposition 9.3 gives $0 \le x_a \le 1$ and every integer complete-signature row sum.

(ii) The joint guarding required there is supplied by Lemmas 5.1 and 6.5, together with the anchor avoidance in Lemma 5.2. In particular, donor occurrences are not separately reinserted.

(iii) The charged base is exactly (10.1); the simultaneous combined charge is (6.50); and the valuation identity after that charge is (10.3).

(iv) The donor-backed row-token construction in Section 5.5 gives nonnegative quotas (6.40), and every complete path obeys the all-factor identity (5.2).

(v) The required two-sided box universality is precisely Lemma 5.3.

Thus there is no additional assumption at the final interface.

Take X, the state indices (ε_g), and $G_{\mathrm{bank}}^{\mathrm{fin}}$ from Proposition 10.2, and define

$$\mathcal{G}_{\mathrm{res}} := G_{\mathrm{fix}} \ \cup \ G_{\mathrm{bank}}^{\mathrm{fin}} \ \cup \{a \in \mathcal{A} : X_a = 1\}.$$

The union is disjoint by that proposition. Since

$$v(B_0) = v(G_{\mathrm{fix}}) + v(G_{\mathrm{bank}}^{0})$$

and the rounded-minus-fractional convention (10.7) gives

$$\sum_{a \in \mathcal{A}} X_a v(a) = \sum_{a \in \mathcal{A}} x_a v(a) + e,$$

the state-replacement identity (10.10) and the exact post-tangent certificate give, at every prime,

$$\begin{aligned} v\left(\prod_{a\in\mathcal{G}_{\mathrm{res}}} a\right) &= v(B_0) + \sum_{a\in\mathcal{A}} x_a v(a) + e - e \\ &= v(Y). \end{aligned}$$

Equality at every prime proves the integer identity

$$\prod_{a\in\mathcal{G}_{\mathrm{res}}} a = Y = \frac{T_n}{D'}.$$

The full central-anchor construction, with the guarded modifications of Section 5.4, gives a set $\mathcal{H}' \subset (n, 2n]$, disjoint from $\mathcal{G}_{\mathrm{res}}$, and

$$\prod_{a\in\mathcal{H}'} a = C_n D'$$

by (5.17). Every factor in $\mathcal{G}_{\mathrm{res}}$ is legal and lies in $(n, M]$: this holds for G_{fix} and the two bank states by their construction and for the flexible set by the definition of $\mathcal{A}$. Hence

$$\mathcal{S} := \mathcal{H}' \cup \mathcal{G}_{\mathrm{res}}$$

consists of distinct integers in $(n, M]$ and satisfies

$$\prod_{a\in\mathcal{S}} a = (C_n D')\frac{T_n}{D'} = C_n T_n = Q(n, M).$$

□

To return to the original formulation, take the complement of the set in Theorem 10.3. Since the product of all integers in $(n, M]$ is $M!/n!$,

$$\prod_{\substack{n<a\le M\\ a\notin\mathcal{S}}} a = \frac{M!/n!}{M!/(n!)^2} = n!.$$

This is a product of distinct integers greater than n, all at most $M = 2n + \lceil cn/\log n\rceil$. Consequently

$$f(n) \le 2n + \left\lceil c\frac{n}{\log n}\right\rceil \qquad (c > C_0)$$

for every sufficiently large n depending on c. Since $c > C_0$ was arbitrary,

$$\limsup_{n\to\infty} \frac{(f(n) - 2n)\log n}{n} \le C_0.$$

Lemma 3.2 gives the reverse liminf. Together they prove Theorem 1.1.

A Numerical verification

The companion numerical verifier is the single human-readable file `numerical_verifier.py`, distributed with this paper. It contains the complete 211-entry rational array used in Lemma 4.1 and certifies precisely the finite claims made there. More specifically, it verifies

- positivity, the allowed cofactor range, and every exact row sum;
- the prime factorization of every cofactor and every finite prime-load capacity inequality;

- the finite–tail overlap inequalities for $201 < p \leq 401$, together with the exact scalar inequality used for the Nagura tail.

The program uses only Python's standard library and `Fraction` arithmetic. It performs no numerical search and makes no floating-point comparison. It therefore checks the finite certificate exactly; it does not replace the analytic tail argument in Lemma 4.2 or any later analytic step of the proof.

References

[ABT99] Richard Arratia, A. D. Barbour, and Simon Tavaré. The Poisson–Dirichlet distribution and the scale-invariant poisson process. *Combinatorics, Probability and Computing*, 8(5):407–416, 1999.

[ACR+26] Boris Alexeev, Evan Conway, Matthieu Rosenfeld, Andrew V. Sutherland, Terence Tao, Markus Uhr, and Kevin Ventullo. Decomposing a factorial into large factors. arXiv preprint, https://arxiv.org/abs/2503.20170, 2026. Version 4, 3 April 2026.

[AS16] Noga Alon and Joel H. Spencer. *The Probabilistic Method*. Wiley Series in Discrete Mathematics and Optimization. John Wiley & Sons, Hoboken, NJ, 4 edition, 2016.

[EGS82] Paul Erdős, Richard K. Guy, and John L. Selfridge. Another property of 239 and some related questions. *Congressus Numerantium*, 34:243–257, 1982.

[HT93] Adolf Hildebrand and Gérald Tenenbaum. Integers without large prime factors. *Journal de Théorie des Nombres de Bordeaux*, 5(2):411–484, 1993.

[LP17] Günter Last and Mathew Penrose. *Lectures on the Poisson Process*. Institute of Mathematical Statistics Textbooks. Cambridge University Press, Cambridge, 2017.

[Mau26] Samuel Mausberg. A thirteen-layer lower bound for Erdős problem #390. Note linked from the Erdős Problems discussion thread, https://www.erdosproblems.com/forum/thread/390, 2026. Dated 2 May 2026.

[Mer74] Franz Mertens. Ein beitrag zur analytischen zahlentheorie. *Journal für die reine und angewandte Mathematik*, 78:46–62, 1874.

[MV07] Hugh L. Montgomery and Robert C. Vaughan. *Multiplicative Number Theory I: Classical Theory*, volume 97 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 2007.

[Nag52] Jitsuro Nagura. On the interval containing at least one prime number. *Proceedings of the Japan Academy*, 28(4):177–181, 1952.

[Sai89] Éric Saias. Sur le nombre des entiers sans grand facteur premier. *Journal of Number Theory*, 32(1):78–99, 1989.

[Tao25] Terence Tao. Comment on Erdős problem #390. Erdős Problems discussion thread, https://www.erdosproblems.com/forum/thread/390, 2025. Posted 30 August 2025.