hexagon:2611.01300v2

ARITHMETIC CONCENTRATION AND RANDOM MATRIX STRUCTURE OF RECIPROCAL SUBSET SUMS

PENGCHENG CHEN

Abstract. We study two arithmetic probability models arising from harmonic subset sums. For a rectangular prime block augmented by a disjoint prime block at height $O(\log p \log_2 p)$, a deletion-stable Fourier estimate yields conditional Gaussian peak profiles and asymptotically independent exposed coordinates on a density-one set of primes. We extend this to joint reciprocal-sum and count coordinates and to all normalized linear projections of polylogarithmically many independent columns conditioned on a common modular label. For the prime block alone, we establish exponential integer-lattice saturation and rank resilience even under modular coincidence, derive an exact Smith identity for scaled numerator minors, and show that the typical gcd is a squarefree product of denominator primes selected by constant rows. A strengthened small-modulus estimate gives a conditional version: after imposing coincidence modulo a moderately growing integer coprime to the denominator product, the gcd is typically the imposed integer times this product, with a Poisson law for the excess prime count in the critical width regime. For fixed column number we also obtain a conditional zeta law in total variation after stripping the imposed factor and the first powers selected by constant rows. We bound factorial moments of the number of large prime divisors and explain the limits of this information for uniform modular coverage. In a separate, uniform distinct-value model we prove concentration of reduced numerator and denominator heights and control the logarithmic mass of small prime factors through recursive Fourier majorants. Neither a conditional Gaussian profile nor a gcd average of fixed order is a pointwise full-coverage theorem.

Contents

1. Introduction and models 2
2. Rectangular blocks and rational reconstruction 2
3. Prime averages, robust mixing, and peak statistics 4
3.1. Exact peak coordinates and conditional independence 5
4. Joint sum and count profiles under modular conditioning 6
5. A pointwise Hamming bound 7
6. Peak moments and arithmetic matrices 8
6.1. Column differences and conditional nondegeneracy 9
7. Random common divisors and Smith invariants 10
7.1. Independent auxiliary columns 12
7.2. Small moduli and large-prime tails 12
7.3. Conditional gcd laws at moderately growing moduli 13
8. Factorial moments of the number of common large primes 15
8.1. The relation to the high-order gcd average 15
9. A uniform distinct-value model: reduced heights and prime factors 16
9.1. The top-level entropy pool 17
9.2. A complementary high-order relation estimate 20
10. Fixed-width conditional zeta law 21
Scope and open questions 22
References 22

Date: 28 November 2026.

2020 *Mathematics Subject Classification.* 11B75, 11D68, 11N37, 60C05, 60B20, 60F05.

Key words and phrases. reciprocal subset sums, Gaussian peaks, conditional laws, Smith normal form, random gcd, Poisson limit.

1. Introduction and models

For a finite set of positive integers D put $R(D) = \sum_{a\in D} 1/a$ and $H_q = R([1, q])$. Intervals of integers are intersected with $\mathbb{Z}$ unless otherwise indicated. Write $\log_j$ for the iterated natural logarithm. For $X \to \infty$ use the dyadic scales $L = \log X$, $U = \log L$, $V = \log U$ and $q_p = \lfloor C \log p \log_2 p \rfloor$ for $X < p \le 2X$, where C is a sufficiently large fixed constant. Set $I_p = [q_p/2, q_p] \cap \mathbb{Z}$. Prime-indexed products and sums are always over primes. The symbol $\|x\|$ denotes distance to the nearest integer; write $e(x) = \exp(2\pi i x)$.

There are two different sources of randomness. In Sections 2–8, the variables are independent fair Bernoulli variables or their independent differences, with weights $w_0(0) = 1/2$ and $w_0(\pm 1) = 1/4$. The numerator of the reciprocal sum of the differences on a block of distinct primes is a one-to-one encoding of the digit vector. We use this to study the gcd of differences of independent such numerators. In Section 9, in contrast, we take a uniform random *distinct rational value* of a large-prime pool, not a uniform representation. This distinction matters for entropy and prime-factor statistics.

The companion manuscript [7] establishes entropy-scale modular coverage and, at the slightly larger scale used here, deletion-stable mixing. We include the necessary spectral input with proof, without counting the coverage theorem again as a new result. The protected-grid multivariate local limit of [8] motivates the sum/count coordinates, but its dense reservoir hypotheses do not hold for the sparse prime block. Our Gaussian argument therefore uses shifted modular Fourier products directly; it yields distribution-function estimates, not exact lattice densities. The high-order gcd criterion of [9] provides the connection with pointwise expansion. We prove a conditional small-modulus gcd law, including a fixed-column zeta limit, and explain why it does not solve that criterion. The classical gcd law for independent uniform integers is discussed by Nymann [2]; the issue here is its validity after arithmetic conditioning and removal of the denominator-prime factors.

The chief new assertions beyond the two source drafts are Theorems 4.1 and 7.5, with the independent-noise extension in Corollary 7.2. The original Gaussian profile, robust rank, Smith identity, typical gcd, factorial moment and reduced-height results are retained with complete proofs. General random integral matrix universality results such as [1] provide context; they are not invoked for our dependent difference rows or conditioned law. We use the prime number theorem with its classical error term, elementary finite Fourier analysis, Hoeffding's and Esseen's inequalities, and Kneser's addition theorem in one auxiliary concentration lemma.

2. Rectangular blocks and rational reconstruction

Fix X, put $L = \log X$, $U = \log L$, $V = \log U$, $Q_0 = CLU$ (unrounded), and use primes $p \in (X, 2X]$. All denominators used below are contained in I_p for large X. In particular, the small $O(1)$ difference between $\log p$ and L creates no endpoint issue. Define the prime sets

$$R_0 = \{r : 8L < r \le 12L\}, \quad S = \{s : CU/14 < s \le CU/13\}, \quad A = \{rs : r \in R_0, s \in S\}.$$

All prime sets in the paper are explicitly prime-indexed. Put $r_0 = |R_0|$, $s_0 = |S|$, $s_- = CU/14$, $s_+ = CU/13$, and $T_S = \prod_{s\in S} s$. Then

$$r_0 \sim 4L/U, \qquad s_0 \sim CU/(182V), \qquad \log T_S = (C/182 + o(1))U.$$

Products in A are distinct by unique factorization, and lie between $4Q_0/7$ and $12Q_0/13$. In particular $|A| \asymp_C L/V$. Write

$$E_F(p, h) = \sum_{a\in F} \left\| ha^{-1}/p \right\|^2,$$

where $\|\cdot\|$ denotes distance to the nearest integer. We use least absolute integer residues for p.

Lemma 2.1 (Uniform weak gap)**.** *For fixed sufficiently large C and every $p \in (X, 2X]$, $h \ne 0$,*

$$E_A(p, h) \gg_C L/(U^2 V).$$

Proof. Call a row r quiet if at least $3s_0/4$ of its phases have distance at most $1/(4s_+)$. For such s, writing u_{rs} for the least absolute residue of $h(rs)^{-1}$, the integers su_{rs} lie in $[-p/4, p/4]$ and are all congruent to hr^{-1}. Thus they coincide with an integer v_r divisible by each quiet s. Write $rv_r = h + k_r p$, $|k_r| = O(L)$. Two quiet rows have at least $s_0/2$ common quiet columns, and their k_r agree modulo each such prime. Their product exceeds $24L + O(1)$ for C sufficiently large, so the two integers k_r agree. All quiet rows therefore divide the same nonzero integer $h + kp$ of size $O(Lp)$, which has at most $(1 + o(1))L/U$ prime divisors in R_0. The remaining $\gg L/U$ rows each contribute at least $s_0/4$ phases of size $\gg 1/U$, giving the claim. □

Lemma 2.2 (Rational reconstruction with lists)**.** *Let $B_0 \geq 1$ be an integer. Denote by $\mathcal{R}_{B_0}$ the set of reduced fractions d/r with $r \in R_0$, $0 < |d| \leq 2r$, $r \nmid d$, and $dr^{-1} \pmod s$ admitting a representative in $[-2B_0, 2B_0]$ for at least $s_0/2$ primes $s \in S$. Let*

$$j_0 = \left\lceil \frac{\log(1000L^2)}{\log s_-} \right\rceil + 1.$$

For sufficiently large fixed C we have $j_0 \leq \lceil s_0/2 \rceil$, and

$$|\mathcal{R}_{B_0}| \leq \frac{\binom{s_0}{j_0}}{\binom{\lceil s_0/2 \rceil}{j_0}} (4B_0 + 1)^{j_0}.$$

In particular, for either $\beta = 1/2$ or $\beta = 3/4$ there exists a fixed $b_0 > 0$ sufficiently small that $B_0 = \lfloor b_0 U^\beta \rfloor$ satisfies

$$|\mathcal{R}_{B_0}| \leq \exp(2\beta U - c_1 U/V). \tag{1}$$

Proof. Choose j_0 qualifying primes and one representative for each. Two reduced fractions having the same record satisfy $s \mid dr' - d'r$ for every recorded prime. The determinant has absolute value at most $576L^2$, less than the product of these primes. Hence it vanishes and the fractions coincide. Each fraction yields at least $\binom{\lceil s_0/2 \rceil}{j_0}$ records. This proves the bound. Now $j_0 = 2U/\log s_- + O(1)$ and $s_0/j_0 \asymp_C 1$. The binomial ratio is at most $\exp(O_C(j_0))$. For $B_0 = \lfloor b_0 U^\beta \rfloor$, choose b_0 so small that the factor per recorded prime (including the binomial ratio) is at most $e^{-2} s_-^\beta$. Then its logarithm times j_0 is at most $2\beta U - c_1 U/V$, since $\log s_- = V + O_C(1)$ and $V = o(U/V)$. This proves (1) with fixed constants. □

Here is the precise comparison of two frequencies used repeatedly. For a pair (p, h) with $1 \leq h < p$, put $\kappa(p, h) = -hp^{-1} \pmod{T_S}$; all the inverses in this expression are legitimate. Let v_r and u_{rs} be the least absolute residues of hr^{-1} and $h(rs)^{-1}$ modulo p. Write

$$rv_r = h + pk_r, \qquad su_{rs} = v_r + pj_{rs}. \tag{2}$$

Thus $|k_r| \leq r/2 + 1$, $|j_{rs}| \leq s\|h(rs)^{-1}/p\| + 1/2$, and

$$k_r + rj_{rs} \equiv \kappa(p, h) \pmod s, \qquad k_r \equiv -hp^{-1} \pmod r. \tag{3}$$

Call a row calm if at least $3s_0/4$ of the columns have $|j_{rs}| \leq B_0$. For two pairs $(p, h), (p', h')$ having the same κ and a common calm row, (3) shows that $d = k_r - k'_r$ satisfies the list condition for at least half the columns. If $r \nmid d$, then $0 < |d| \leq r + 2 \leq 2r$, and $d/r \in \mathcal{R}_{B_0}$.

Proposition 2.3 (Pointwise small spectrum and robust prime-average gap)**.** *There exist fixed positive constants c_0, c_1, δ such that the following hold.*

(1) *For every $p \in (X, 2X]$, if $\tau_1 = c_0 L/(UV)$, then*

$$\#\{h \neq 0 : E_A(p, h) \leq \tau_1\} \leq T_S = \exp(O_C(U)).$$

(2) *Put $\tau_2 = c_1 L/(\sqrt{U}V)$. Apart from $O(Xe^{-cL/V})$ primes in $(X, 2X]$, every nonzero h and every $D \subseteq A$ with $|D| \leq \delta|A|$ satisfy*

$$E_{A \setminus D}(p, h) > \tau_2.$$

Proof. In either case the number of entries with $|j_{rs}| > B_0$ is at most $|D| + 4s_+^2 E_{A\setminus D}/B_0^2$, since for such an entry the phase distance is at least $B_0/(2s_+)$. Choose the threshold constants and δ so that this bound is at most $|A|/32$. It follows that at least $7r_0/8$ rows are calm. For (1) take $D = \varnothing$, $B_0 = \lfloor b_0\sqrt{U}\rfloor$. By (1), only $o(r_0)$ rows admit any fraction of the exceptional list. Two frequencies with the same κ then agree modulo r for at least $(3/4 - o(1))r_0$ rows. Their prime product exceeds p, forcing $h = h'$. This gives the claimed injection into the T_S residue classes.

For (2) take $B_0 = \lfloor b_0 U^{3/4}\rfloor$. For fixed κ and row r, choose one calm pair as representative if one exists. Each other calm pair has either the same k_r residue or corresponds, by subtraction from this representative, to a member d/r of the list. Therefore there are at most $1 + D_r$ possible k_r residues for calm pairs, where D_r is the number of list fractions of denominator r. We have $\sum_r D_r = |\mathcal{R}_{B_0}|$. Choose $\ell = \lceil \log(4X^2)/\log(8L)\rceil + 1$ calm rows; this is possible since $\ell = (2+O(1/U))L/U$. Equality of their residues for two pairs implies that each of those primes divides $hp' - h'p$. Their product exceeds $4X^2$, so the reduced fractions $h/p, h'/p'$ coincide. Thus for a specified row set the number of pairs is at most the product of the corresponding $1 + D_r$. Summing over row sets and κ, and using AM–GM, bounds the total number of bad pairs by

$$T_S 2^{r_0}\left(1 + \frac{|\mathcal{R}_{B_0}|}{\ell}\right)^{\ell}.$$

By (1), the last logarithm per factor is at most $U/2 - c'U/V$, since $\log \ell = U - V + O(1)$ and $V = o(U/V)$. As $\log T_S + r_0 \log 2 = O_C(U + L/U) = o(L/V)$, the number of pairs is $O(Xe^{-cL/V})$. No union over deletion sets is needed: the property counted is existence of enough calm rows. The claim follows. □

3. Prime averages, robust mixing, and peak statistics

Use the prime block

$$B = \{b \text{ prime} : 4Q_0/5 < b < 9Q_0/10\}, \quad n = |B| = (C/10 + o(1))L.$$

It is disjoint from A. Write $F = A \cup B$. Its total reciprocal mass is $O_C(1/U)$ and $|F| \asymp_C L$. For any set D of invertible denominators and independent Bernoulli variables ϵ_a of parameter $1/2$, define

$$\phi_{p,D}(h) = \prod_{a\in D} (1 + \mathrm{e}^{2\pi i h a^{-1}/p})/2.$$

The independent difference variables $C_a = \epsilon_a - \epsilon'_a$ take $-1, 0, 1$ with probabilities $1/4, 1/2, 1/4$, and their transform is $|\phi_{p,D}(h)|^2$. We use both $|\cos \pi x| \le e^{-2\|x\|^2}$ and $\cos^2 \pi x \le e^{-\sin^2 \pi x}$.

Lemma 3.1 (Poisson majorant on average)**.** *For C sufficiently large and*

$$G_{p,B}(h) = \exp\left(-\tfrac{1}{8}\sum_{b\in B} \sin^2(\pi h b^{-1}/p)\right),$$

we have $\sum_{X<p\le 2X} \sum_{h \bmod p} G_{p,B}(h) \ll_C XU$.

Proof. Take independent $K_b = N_b^+ - N_b^-$ with $N_b^{\pm} \sim \mathrm{Pois}(1/32)$. Then $G_{p,B}$ is the transform of $\sum K_b/b$ modulo p. Truncate at $\sum |K_b| \le n$; the probability of the complementary event is $\exp(-cn)$ by domination by a Poisson variable of mean $n/16$. On the truncated set the rational sum is zero only if all coefficients vanish: multiplication by $P = \prod_{b\in B} b$ and reduction modulo each b proves this since $n < \min B$. The probability of the all-zero vector is $w(0)^n$ for a fixed $w(0) < 1$. With C large both errors are $O(X^{-3})$. Each nonzero truncated numerator with denominator P has logarithmic size $O_C(LU)$ and hence at most $O_C(U)$ prime divisors in $(X, 2X]$. Orthogonality says the sum over h is p times the zero-residue probability. Average this identity over p, bounding each nonzero numerator's weighted prime count by $O(XU)$. The two exceptional probabilities contribute $O(X^2X^{-3})$. This proves the result. □

Theorem 3.2 (Deletion-stable mixing)**.** *With $\tau = \tau_2$ from Proposition 2.3 (decreasing its constant if necessary), there is a set of good primes in $(X, 2X]$, with complement $O(Xe^{-c\tau})$, such that for every good prime and every $D \subseteq F$ satisfying $|D \cap A| \le \delta|A|$ and $|D \cap B| \le \tau$,*

$$\sum_{h\neq 0} |\phi_{p,F\setminus D}(h)| \le e^{-\tau}. \tag{4}$$

For either Bernoulli or difference variables the residual sum modulo p therefore has point probabilities $p^{-1}(1 + O(e^{-\tau}))$, uniformly in D and the residue.

Proof. Restrict to primes satisfying Proposition 2.3(2) and $\sum_{h\neq 0} G_{p,B}(h) \le e^{\tau/2}$. The omitted primes number $O(Xe^{-cL/V} + XUe^{-\tau/2}) = O(Xe^{-c'\tau})$ by Lemma 3.1 and Markov's inequality. For a remaining prime, the A part after deletion contributes at most $e^{-2\tau}$ to each nonzero coefficient. On the B part use $|\cos \pi x| \le e^{-\sin^2(\pi x)/2} \le e^{-\sin^2(\pi x)/8}$. Deleting at most τ factors can increase the Poisson majorant by at most $e^{\tau/8}$. Consequently the left side of (4) is at most $e^{-2\tau+\tau/8+\tau/2} \le e^{-\tau}$. Squaring coefficients only improves the bound. Fourier inversion gives the conclusion. □

The same Fourier bound holds after convolution with the independent choices on $I_p \setminus F$, since Fourier magnitudes cannot increase.

3.1. **Exact peak coordinates and conditional independence.** For a set D of positive denominators with total reciprocal mass less than $1/4$, let $T_D = \operatorname{lcm}(D)$, and take $p > \max D$. For the difference variables define $R_D = \sum_{a\in D} C_a/a$ and $Y_D = \bar{p}R_D \pmod 1$, where $p\bar{p} \equiv 1 \pmod{T_D}$. There is an exact identity on the circle

$$Y_D = \frac{J + R_D}{p} \pmod 1, \qquad J \equiv -\sum_{a\in D} C_a a^{-1} \pmod p. \tag{5}$$

Indeed, with $N = T_D R_D \in \mathbb{Z}$, write $p\bar{p} = 1 + kT_D$. Then $pY_D \equiv R_D \pmod 1$ and $kN \equiv -NT_D^{-1} \pmod p$. The interval of possible displacements R_D/p is strictly smaller than half the peak spacing. For $F = A \cup B$ or B alone, the mass condition holds for large X. For difference variables all finite-field Fourier coefficients are nonnegative, and thus the largest peak mass equals the central one, $\pi_p(0) = \Pr(J = 0)$.

Corollary 3.3 (Conditional coordinate statistics)**.** *For the good primes of Theorem 3.2, take $F = A \cup B$, and let J be its peak label. For any admissible deletion set D as in that theorem and any configuration c_a, $a \in D$,*

$$\mathbb{P}(C_a = c_a \ (a \in D) \mid J = j) = (1 + O(e^{-\tau})) \prod_{a\in D} \mathbb{P}(C_a = c_a)$$

uniformly in j, D and the configuration. The same holds for Bernoulli coordinates with the analogous label.

Proof. For each configuration the residual independent sum is equidistributed with the relative error in Theorem 3.2; the total label is also equidistributed. Bayes' rule gives the formula, including for configurations of very small probability. In particular all bounded centered products of these specified coordinates have conditional expectation $O(e^{-\tau})$. □

Theorem 3.4 (Gaussian peak profiles)**.** *On the same good primes let $R_F = \sum_{a\in F} C_a/a$, $\sigma^2 = \frac{1}{2}\sum_{a\in F} a^{-2} \asymp_C (LU^2)^{-1}$, and Φ be the standard normal distribution function. Then*

$$\sup_{j,x} |\mathbb{P}(R_F/\sigma \le x \mid J = j) - \Phi(x)| \ll_C \tau^{-1/2}.$$

Consequently $\mathbb{P}(J = j, R_F/\sigma \le x) = p^{-1}(\Phi(x) + O_C(\tau^{-1/2}))$ uniformly. This is a distribution-function approximation, not a local density theorem for the discrete measure.

Proof. Use the conditional characteristic function $\varphi_j(t) = \mathbb{E}(e^{itR_F/\sigma} \mid J = j)$. Finite Fourier inversion expresses $p\mathbb{P}(J = j)\varphi_j(t)$ as a sum over $h \bmod p$ of factors of the form

$$\mathrm{e}^{-2\pi ihj/p} \prod_{a\in F} \cos^2 \pi \left(-ha^{-1}/p + t/(2\pi a\sigma)\right).$$

The $h = 0$ term is the unconditional characteristic function $\varphi(t)$. For $h \neq 0$ use $\|x + y\|^2 \geq \frac{1}{2}\|x\|^2 - |y|^2$ and $\cos^2(\pi x) \leq e^{-4\|x\|^2}$. The product is bounded by

$$\exp(2t^2/\pi^2) \exp(-2E_A(p, h))G_{p,B}(h),$$

since $2\|x\|^2 \geq (1/8)\sin^2(\pi x)$ and $\sum_{a\in F}(t/(2\pi a\sigma))^2 = t^2/(2\pi^2)$. For good primes, $E_A(p, h) > \tau$ and $\sum_{h\neq 0} G_{p,B}(h) \leq e^{\tau/2}$. Uniformly for $|t| \leq \sqrt{\tau}$, the error after normalization by $p\mathbb{P}(J = j)$ is $O(e^{-\tau})$.

As $a \asymp_C LU$ and $|F| \asymp_C L$, Taylor expansion of the cosine yields

$$\varphi(t) = \prod_{a\in F} \cos^2(t/(2a\sigma)), \qquad |\varphi(t) - e^{-t^2/2}| \ll_C (t^4/L)e^{-t^2/2}$$

for $|t| \leq \sqrt{\tau}$: the logarithmic correction is nonpositive and of size $O_C(t^4/L)$, even when this size is not small. At the origin the uniform exponential error cannot simply be divided by $|t|$. Both conditional and unconditional random variables have absolute value $O_C(\sqrt{L})$, so $|\varphi_j(t) - \varphi(t)| \ll_C \sqrt{L}|t|$ as well. Using the smaller of this bound and $O(e^{-\tau})$, its integral divided by $|t|$ up to $\sqrt{\tau}$ is $O_C(e^{-\tau}(1 + \tau + \log L))$. The Gaussian comparison term integrates to $O_C(1/L)$. Esseen's smoothing inequality with cutoff $\sqrt{\tau}$ proves the claim. □

4. Joint sum and count profiles under modular conditioning

The real peak variable can be augmented by a count coordinate, without assuming that the underlying exact rational grid is aperiodic. For differences the count is the difference of the two subset cardinalities. Fix a good prime as in Theorem 3.2; the proof there also gives the energy bound on A and the Poisson-majorant bound on B. For an admissible D write $F_D = F \setminus D$ and

$$V_D = \sum_{a\in F_D} C_a \begin{pmatrix} Q_0/a \\ 1 \end{pmatrix}, \qquad \Sigma_D = \mathrm{Cov}(V_D) = \tfrac{1}{2} \sum_{a\in F_D} \begin{pmatrix} Q_0/a \\ 1 \end{pmatrix} (Q_0/a, 1), \qquad W_D = \Sigma_D^{-1/2} V_D.$$

The covariance is uniformly nonsingular: $c_C LI_2 \leq \Sigma_D \leq C_C LI_2$. For the lower bound, take primes in two fixed separated subintervals of $(4Q_0/5, 9Q_0/10)$, each containing $\gg_C L$ surviving primes by the prime number theorem and $|D \cap B| \leq \tau = o(L)$. The two sets of directions $(Q_0/b, 1)$ are separated by a fixed positive angle. For every unit vector, at least one of the subintervals (which can be chosen sufficiently short) has scalar products bounded below in absolute value. Summing squares proves the bound. The upper bound follows from $|F| = O_C(L)$ and $a \asymp_C Q_0$.

Theorem 4.1 (Joint conditional Gaussian projections)**.** *Fix $d > 0$. Take $1 \leq m \leq U^d$ independent difference columns on F and their respective peak labels $J_1, \ldots, J_m$. In column i fix an admissible deletion set D_i before sampling, expose its coordinates and prescribe any configuration on it; define $W_i = W_{D_i}$ from the unexposed variables in that column. Conditional either on any prescribed vector of labels or on the event $J_1 = \cdots = J_m$, and on the prescribed configurations, uniformly over all such choices and vectors $u_i \in \mathbb{R}^2$ with $\sum_i |u_i|^2 = 1$,*

$$\sup_x \left| \mathbb{P}\left(\sum_{i=1}^m u_i^\top W_i \leq x \;\middle|\; \textit{the stated conditions} \right) - \Phi(x) \right| \ll_{C,d} \tau^{-1/2}.$$

Moreover the joint law of the exposed configurations, conditional on either label condition alone, is their unconditional product law times $1 + O(me^{-\tau})$, uniformly even for individual configurations of very small probability. For fixed m this gives joint weak convergence to the standard $2m$-dimensional Gaussian for the normalized residual vectors, uniformly in the label

conditions in the sense of convergence of all linear projections. No growing-dimensional density or total variation conclusion is asserted.

Proof. First consider one residual column. Write $v_a = (Q_0/a, 1)^\top$ and $\lambda_a = u^\top \Sigma_D^{-1/2} v_a$ for a unit vector u. We have $\sum \lambda_a^2 = 2$ and $\max|\lambda_a| \ll_C L^{-1/2}$. For a residual label j, Fourier inversion of the label together with the characteristic function of $u^\top W_D$ gives products

$$\prod_{a \in F_D} \cos^2 \pi(-ha^{-1}/p + t\lambda_a/(2\pi)).$$

The $h = 0$ term is its unconditional characteristic function $\varphi_{D,u}(t)$. Using $\|x+y\|^2 \geq \|x\|^2/2 - |y|^2$, the sum of the absolute nonzero terms is bounded by

$$\exp(2t^2/\pi^2 - 2\tau)\exp(\tau/8)\sum_{h\neq 0} G_{p,B}(h) \leq \exp(2t^2/\pi^2 - 11\tau/8).$$

Here the residual A retains energy greater than τ, and deletion in B costs at most $e^{\tau/8}$. Consequently, after division by $p\mathbb{P}(J_{F_D} = j) = 1 + O(e^{-\tau})$, the conditional characteristic function differs from the unconditional one by $O(e^{-\tau})$ for $|t| \leq \sqrt{\tau}$. Prescribed exposed variables simply shift the desired residual label.

For prescribed full labels the residual columns are independent conditional on these exposures and labels. Put $r_i = |u_i|$, use unit directions u_i/r_i when $r_i > 0$, and apply the preceding bound at $r_i t$. Telescoping the product (all characteristic functions have modulus at most one) gives an error $O(me^{-\tau})$. The corresponding unconditional product is $\prod_{i,a} \cos^2(t\lambda_{ia}/2)$, where $\sum_{i,a} \lambda_{ia}^2 = 2$ and $\max_{i,a}|\lambda_{ia}| \ll_C L^{-1/2}$. Since $\tau = o(L)$, Taylor expansion of $\log\cos$ shows, uniformly for $|t| \leq \sqrt{\tau}$, that the logarithm equals $-t^2/2$ plus a nonpositive error of magnitude $O_C(t^4/L)$. Thus its difference from $e^{-t^2/2}$ is $O_C(t^4 e^{-t^2/2}/L)$. All variables in the projection have absolute total value $O_C(\sqrt{mL})$ by Cauchy–Schwarz, so the conditional–unconditional difference is also $O_C(\sqrt{mL}|t|)$. Esseen's smoothing inequality, as in Theorem 3.4, bounds the error by

$$O_C\left(\tau^{-1/2} + L^{-1} + me^{-\tau}(1 + \tau + \log(mL))\right) \ll_{C,d} \tau^{-1/2}.$$

The bound is uniform over every vector of labels. Under equality of all labels the conditional law is a mixture of those with prescribed equal labels; the uniform estimate survives averaging.

Finally, for a given vector of exposures each individual residual-label probability is $p^{-1}(1 + O(e^{-\tau}))$ uniformly. The probability of any prescribed full label vector is therefore $p^{-m}(1 + O(me^{-\tau}))$, and of equality of labels is $p^{1-m}(1 + O(me^{-\tau}))$, with or without the exposures. Bayes' rule gives the asserted relative product law. Fixed-dimensional convergence follows from the Cramér–Wold theorem. □

The proof uses the sum/count covariance mechanism familiar from [8] but does not use the protected-grid local limit theorem there. It supplies a genuinely conditional multivariate Gaussian statement for sparse reciprocal blocks without making an unjustified assertion of exact count-lattice uniformity.

5. A pointwise Hamming bound

For the difference law on B, let $\rho_p(B)$ denote its largest point mass.

Lemma 5.1 (Weighted Hamming bound)**.** *Let B' be a set of distinct primes in I_p and write $n' = |B'|$. For any integer modulus $q \geq 2$ coprime to their product, define $t = \min(n', \lfloor \log q/\log q_p \rfloor)$, $s = \lfloor t/2 \rfloor$. For the independent three-valued differences C_b (or independent Bernoulli variables) the maximal point probability of $\sum_{b\in B'} C_b b^{-1} \pmod q$ is at most*

$$\beta(q) = \min\left(2^{-t}, \left[\sum_{j=0}^{s}\binom{n'}{j}\right]^{-1}\right). \tag{6}$$

Proof. Two distinct input vectors differing in at most t coordinates have a nonzero rational difference (reduce its numerator modulo one of the distinct prime denominators). The common denominator of the difference is at most $q_p^t \le q$. The real difference has magnitude at most $2n'/\min B' = O_C(1/U) < 1$. Thus its nonzero numerator has absolute value less than q, ruling out equality modulo q. Every fibre is a code of minimum Hamming distance $> t$. Fixing all but t coordinates gives the first bound. For the second, each coordinate weight is at most $1/2$, so the combined weight of other choices divided by the original coordinate weight is at least one. A Hamming ball of radius s therefore has product mass at least the mass of its centre times $\sum_{j\le s}\binom{n'}{j}$. Disjoint balls around fibre elements prove the estimate. The case $t = 0$ is covered by the trivial bound. □

For $B' = B$ and $q = p$, Stirling's estimate gives $\rho_p(B) \le \exp(-(1/2 - o(1))LV/U)$. This is much smaller than a polynomial in L^{-1} but far larger than the needed bound $p^{-1}e^{o(\tau_1)}$.

6. Peak moments and arithmetic matrices

Use the prime block B alone and put $P = \prod_b b$. Its integer numerators $N = \sum_b C_b P/b$ are injective as functions of the three-valued input, since reducing an equality modulo b forces equality of its b coordinate. Let $\pi_p(j)$ be the peak masses in (5), for this B, and $S_m(p) = \sum_j \pi_p(j)^m$. Set $Y_0 = \lceil KP/(pU)\rceil$ with K sufficiently large and fixed. Define the probability on $\mathbb{Z}/P\mathbb{Z}$

$$F_p(z) = \prod_{b\in B} w_0(pz(P/b)^{-1} \bmod b), \qquad w_0(0) = 1/2, \quad w_0(\pm 1) = 1/4,$$

where w_0 is zero elsewhere and periodized in each local modulus. For shifts $\mathbf{h} = (h_1, \dots, h_m)$ set

$$\mathcal{K}_m(\mathbf{h}) = \sum_{z \bmod P} \prod_{i=1}^{m} F_p(z + h_i).$$

For any $m \ge 2$, the exact peak separation identity is

$$S_m(p) = \sum_{|h_i|_P \le Y_0,\, 2\le i\le m} \mathcal{K}_m(0, h_2, \dots, h_m). \tag{7}$$

Here $|h|_P$ is the least absolute cyclic representative. To prove it, the right side is the probability that m independent CRT points all lie within cyclic distance Y_0 of the first. By (5), points from one peak have diameter $O(P/(pU))$, while points from different peaks have separation at least $P/p - O(P/(pU)) > Y_0$. Thus the condition is exactly equality of all peak labels.

For $\alpha_b = p(P/b)^{-1} \pmod b$ and $d_b = m - \max_v \#\{i : \alpha_b h_i = v \pmod b\}$, the local complete correlation is $\mathcal{K}_{b,m} = \sum_{x \bmod b} \prod_i w_0(x + \alpha_b h_i)$, with $\mathcal{K}_m = \prod_b \mathcal{K}_{b,m}$. It obeys

$$\frac{\mathcal{K}_{b,m}(\mathbf{h})}{2^{-m} + 2\cdot 4^{-m}} \le 2^{-d_b/2}. \tag{8}$$

If $d_b = 0$ this is equality. If $d_b \ge 2$, at most two shifts can give a nonzero product (when two distinct local residues occur), and each product is at most 2^{-m-d_b}. For $d_b = 1$ there are two residue classes of multiplicities $m-1$ and 1. For distance ± 1 the numerator is $2^{-(m+1)} + 2^{-(2m-1)}$, whose ratio to $2^{-m} + 2\cdot 4^{-m}$ is at most $2/3$ for $m \ge 2$. For distance ± 2 the numerator is 4^{-m}, and for other differences it is zero. We use $b > 5$ throughout. In particular for mutually distinct shifts in $[-Y_0, Y_0]$, the prime factors of their pairwise differences show that at least $\log(P/(2Y_0))/\log q_p$ rows distinguish each pair. At most md_b unordered pairs can differ in a row, since each such pair has an index outside a maximal residue class. Thus

$$\frac{\mathcal{K}_m(\mathbf{h})}{\mathcal{K}_m(\mathbf{0})} \le \exp\left(-\tfrac{\log 2}{4}(m-1)\frac{\log(P/(2Y_0))}{\log q_p}\right).$$

Repeated CRT points have relatively negligible contribution to (7) when $m = O((UV)^2)$: for any pair of coincident points and all other points in that peak, the total mass is at most $2^{-n}S_{m-1}(p)$. Since $S_{m-1} \le pS_m$ by Hölder and $S_m \ge p^{1-m}$, the union bound gives a relative

error at most $\binom{m}{2}p2^{-n} = o(1)$ for large C. This is a relative bound, rather than an absolute estimate divided by an exponentially smaller target.

6.1. **Column differences and conditional nondegeneracy.** In what follows $3 \le m \le \lceil (UV)^2 \rceil + 2$, and $k = m - 1$. Let (c_{bi}) be an $n \times m$ matrix with independent entries of law w_0, define $N_i = \sum_b c_{bi} P/b$, and let M be the difference matrix with rows

$$\mathbf{v}_b = (c_{b2} - c_{b1}, \dots, c_{bm} - c_{b1}) \in \mathbb{Z}^k.$$

The vector $\mathbf{A} = \sum_b (P/b)\mathbf{v}_b$ consists of the differences $N_i - N_1$, $i \ge 2$. For any $q \ge 2$ let $\mathcal{E}_q = \{\mathbf{A} \equiv 0 \pmod q\}$ and $S_j(q) = \sum_x \mu_q(x)^j$, where μ_q is the law of $N \pmod q$. In particular $\mathbb{P}(\mathcal{E}_q) = S_m(q)$; $S_m(p)$ agrees with the peak moment above (a permutation of residues). For any distribution on q points (including zero weights), Hölder gives

$$S_{m-2}(q)/S_m(q) \le q^2 \tag{9}$$

for $m \ge 3$: $S_{m-2} \le q^{2/m} S_m^{(m-2)/m}$ and $S_m \ge q^{1-m}$, so their ratio is at most $q^{2/m} q^{2(m-1)/m} = q^2$. This permits conditioning on two columns without paying the much larger reciprocal of the entire same-peak probability.

Finite lemma 1 (Finite conditional witness estimate)**.** *Let $n \ge k = m - 1 \ge 2$. Take independent entries of the three-point law w_0 in an $n \times m$ matrix, and arbitrary fixed integer weights a_b in each column numerator $N_i = \sum_b a_b c_{bi}$. Write M for the difference matrix as above and $\mathcal{E}_q$ for equality of the numerators modulo $q \ge 2$. Then, with $H = (2k)^k$,*

$$\mathbb{P}(\Lambda(M) \ne \mathbb{Z}^k \mid \mathcal{E}_q) \le H \binom{m}{2} q^2 \sum_{j=0}^{k-1} \binom{n}{j} 2^{-(n-j)}. \tag{10}$$

The analogous unconditional estimate omits q^2.

Proof. All maximal minors of M are at most H in absolute value by Hadamard (each row has entries of magnitude at most two). Nonsaturation implies rank deficiency modulo a prime $\ell \le H$: if the rational rank is deficient take $\ell = 2$, otherwise take a prime divisor of the index of the row lattice, which divides each nonzero maximal minor. Fix such a prime. On the $(m-1)$-dimensional sum-zero subspace of $\mathbb{F}_\ell^m$, rank deficiency is the existence of a nonzero vector annihilating the row forms. A maximal independent set of those forms uses at most $k - 1$ rows. Fix a choice of this witness set E of size j, and, from its entries, select a nonzero kernel vector by a fixed deterministic rule if one exists. For this witness, the vector annihilates all rows on the deficiency event when the selected forms span the row forms. It has at least two nonzero coordinates because their sum is zero.

Union bound over the $\binom{m}{2}$ possible pairs of these coordinates. For a fixed pair, expose the other $m - 2$ columns and the two selected columns in the witness rows. The chosen vector is fixed by this exposure. If either selected coefficient is zero the case contributes nothing. Otherwise on each of the $n - j$ remaining rows the independent two entries have linear-combination maximum atom at most $1/2$. To see this, for nonzero multiples of the two laws apply Cauchy–Schwarz to the translated convolution sum: each squared-mass sum is $3/8$ for odd ℓ and $1/2$ for $\ell = 2$ (the two nonzero digits then coincide). Thus $1/2$ is uniform, and $3/8$ suffices for odd ℓ. The probability of annihilating all the remaining rows is at most $2^{-(n-j)}$ conditionally on the exposure.

The event $\mathcal{E}_q$ implies equality already among the other $m-2$ columns; the latter is measurable with respect to the exposure, and has probability $S_{m-2}(q)$. Multiplying by the conditional row bound before averaging, then summing over at most H primes, witness sets and column pairs, bounds the joint probability by the right side of (10) with q^2 replaced by $S_{m-2}(q)$. Divide by $S_m(q) > 0$ and apply (9). Without the same-label event the argument omits the factor $S_{m-2}(q)$ and the division. No conditioning on the actual kernel of the full random matrix is used. □

Theorem 6.1 (Robust rank and integer lattice saturation)**.** *For C sufficiently large and all $q \le 2X$ with $q \ge 2$,*

$$\mathbb{P}(\Lambda(M) \neq \mathbb{Z}^k \mid \mathcal{E}_q) \le \exp(-n \log 2 + 2 \log q + o(L)),$$

where $\Lambda(M)$ is the lattice generated by the rows. The same unconditional bound holds without $2 \log q$. In particular the conditional probability of saturation tends to one with exponential error for all target primes $p \in (X, 2X]$. Moreover, for any fixed $\varepsilon \in (0, 1/8)$, after choosing C sufficiently large depending on ε, the following holds under $\mathcal{E}_p$: with conditional probability at least $1 - e^{-c_\varepsilon L}$,

$$(5/8 - \varepsilon)n \le \min_{\lambda \neq 0,\ \sum_i \lambda_i = 0} \#\{b : \sum_i \lambda_i c_{bi} \neq 0\} \le (5/8 + \varepsilon)n,$$

where λ ranges over $\mathbb{F}_p^m$. Thus removing fewer than $(5/8 - \varepsilon)n$ rows preserves full column rank of the difference matrix M over $\mathbb{F}_p$, hence also over $\mathbb{Q}$. Fixed C does not here permit arbitrary $\varepsilon \to 0$.

Proof. For the saturation assertion apply Finite lemma 1 with weights $a_b = P/b$. Here $\log((2k)^k) = o(L)$, $k = o(n)$, and

$$\binom{m}{2} \sum_{j=0}^{k-1} \binom{n}{j} 2^j = \exp(o(L)),$$

by $m = O((UV)^2)$ and $n \asymp_C L$. Factoring out 2^{-n} gives the claimed bound and its unconditional counterpart. We will also use the sharper Hadamard bound $H_k = (2\sqrt{k})^k$ for the size of a maximal minor. For the support assertion, if a nonzero sum-zero vector annihilates some collection of rows, those row forms restricted to the sum-zero space span a space of dimension at most $m - 2$. As in the finite lemma, choose a basis witness set of at most $m - 2$ rows. The deterministic nonzero vector in its kernel annihilates all those rows. The same two-column exposure works, with the sharper $3/8$ maximal atom at odd primes. The number of witness sets and column pairs is again $\exp(o(L))$.

For robust rank, use the same witnesses modulo p, and the bound $3/8$ on the conditional zero probability at an unfixed row. Hoeffding gives $e^{-c\varepsilon^2 n}$ for more than $(3/8 + \varepsilon)n$ zero rows (the $O(m)$ witness rows are negligible). The union bound and (9) give conditional error $\exp(2L - c\varepsilon^2 n + o(L))$. The upper bound for the minimum support follows by considering the difference of the first two columns: their support is binomial of parameter $5/8$ without conditioning, and the same two-column argument applies to its upper tail. Increase C to get exponential error. This proves the theorem. □

This is row-deletion resilience for the dependent difference matrix under modular coincidence. The Rademacher entry-change resilience theorem of Ferber, Luh and McKinley [4] is a different model and is not used here.

For completeness, typical same-peak matrices have total defect $\sum_b (m - \max_v \#\{i : c_{bi} = v\})$ within εnm of $nm/2$ for $m \to \infty$. The expectation differs by $O(n\sqrt{m})$, since each row has multinomial counts of probabilities $(1/4, 1/2, 1/4)$. Changing one entry changes total defect by at most one. Bounded differences gives failure probability $e^{-c\varepsilon^2 nm}$ unconditionally, and dividing by $S_m(p) \ge p^{1-m}$ gives conditional failure $e^{-c'\varepsilon^2 nm}$ with C sufficiently large. Equation (8) (or its bound before division) thus implies $\mathcal{K}_m(\mathbf{h}) \le 2^{n-(3/2-\varepsilon)nm}$ on shifts accounting for almost all peak-moment mass. Neither this nor robust rank estimates the number of such shifts to the required exponential precision.

7. Random common divisors and Smith invariants

Keep the setting of Section 6, with $k = m - 1$, $P = \prod_{b \in B} b$ and $\mathbf{A} = (N_2 - N_1, \dots, N_m - N_1)$. Write $G = \gcd(A_1, \dots, A_k) \ge 0$ (set $G = 0$ for the zero vector). If M has full column rank

then $\mathbf{A} \neq 0$: otherwise all N_i coincide, injectivity of the rational sums forces all columns equal, contradicting full rank. For $E \subset B$ with $|E| = k-1$ put $P_E = \prod_{b\in E} b$ and

$$D_E = \frac{\det(M_E, \mathbf{A})}{P_E} = \sum_{b\notin E} \frac{P/P_E}{b} \det(M_E, \mathbf{v}_b), \qquad \mathfrak{G}(M) = \gcd_{D_E\neq 0} |D_E|.$$

The equality follows by linearity of the last row, and shows $D_E \in \mathbb{Z}$. On the full-rank event there is a nonzero D_E because a row basis spans the entire rational space and $\mathbf{A} \neq 0$. The signs associated with the row order do not affect any gcd.

Finite lemma 2 (Scaled determinant gcd on a saturated row lattice)**.** *Let B be a finite set of distinct primes $b > 2$, $P = \prod_{b\in B} b$, and $\mathbf{v}_b \in \mathbb{Z}^k$ with entries in $[-2, 2]$. Suppose these rows span $\mathbb{Z}^k$ as an integer lattice and $\mathbf{A} = \sum_b (P/b)\mathbf{v}_b \neq 0$. For each $(k-1)$-element row set E let D_E be the scaled determinant defined above. Then all D_E are integers, some $D_E \neq 0$, and*

$$\gcd_E |D_E| = \gcd_{1\le i\le k} |A_i|.$$

Proof. The last-row expansion gives the integer expression for D_E above, and spanning over $\mathbb{Q}$ ensures at least one determinant is nonzero. Denote by G the coordinate gcd of $\mathbf{A}$, and by G' the gcd of the unscaled determinants. Expansion along the last row shows $G \mid G'$. Conversely, as the rows span $\mathbb{Z}^k$, each coordinate unit vector is an integer combination of them. For each i, the determinant of the $k-1$ coordinate unit vectors other than the ith, with last row $\mathbf{A}$, equals $\pm A_i$. Multilinearity expresses this as an integer combination of the unscaled determinants (repeated rows have zero determinant). Thus $G' \mid G$, so $G' = G$. This step uses only integer multilinearity, not Smith normal form.

For any prime $\ell \nmid P$, all scaling denominators P_E are units at ℓ, so the scaled gcd has the same ℓ-valuation as G'. For $b \mid P$, first suppose $\mathbf{v}_b \neq 0$. Its reduction modulo b is nonzero, since its entries have magnitude at most two. The equality $\mathbf{A} \equiv (P/b)\mathbf{v}_b \pmod b$ shows $b \nmid G$. The row-span condition implies full span modulo b; extend the nonzero row $\mathbf{v}_b$ to a basis using other rows E. In the integer expansion of D_E modulo b, the only surviving term is the one indexed by b, with nonzero coefficient and determinant. Hence b does not divide the scaled gcd either. If $\mathbf{v}_b = 0$, any nonzero determinant excludes row b; for those determinants $b \nmid P_E$ and scaling preserves their b-valuations. All valuations of the two positive gcds now agree. This proves the finite identity. □

Theorem 7.1 (Smith identity for the reciprocal numerators)**.** *On the full-rank event, let $d_1 \mid \cdots \mid d_k$ be the Smith factors of the row lattice and $\delta_0 = \prod_i d_i$. In suitable unimodular coordinates write $A'_i = d_i a_i$. Then*

$$\gcd_E |\det(M_E, \mathbf{A})| = \delta_0 \gcd_i |a_i|.$$

For every prime $\ell \nmid P$, this implies $v_\ell(\mathfrak{G}) = v_\ell(\delta_0) + \min_i v_\ell(a_i)$ (with the usual infinite valuation for zero coordinates). In particular, if $\ell > H_k = (2\sqrt{k})^k$ and $\ell \nmid P$, then $v_\ell(\mathfrak{G}) = v_\ell(G)$. If the row lattice is saturated, then $\mathfrak{G}(M) = G$ as integers, even at primes dividing P.

Proof. The ideal generated by the determinants with the last row fixed at $\mathbf{A}$ is unchanged by taking all $(k-1)$-tuples from the row lattice: expand integer combinations by multilinearity. It is unchanged by a unimodular coordinate transformation. In the Smith basis $d_i e_i$, taking all but the ith basis vector produces determinant $\pm(\delta_0/d_i)A'_i = \pm\delta_0 a_i$, giving the stated gcd. Division by P_E does not change valuations away from P. Also δ_0 divides any nonzero maximal minor, and for $\ell > H_k$ all Smith factors are units modulo ℓ^j for every j. The unimodular coordinate change preserves G.

For the saturated case Finite lemma 2 gives the full integer identity, including the denominator primes, without appealing to the Smith coordinate calculation. □

Theorem 6.1 shows that saturation holds with high probability even conditional on $\mathcal{E}_p$. For the target large primes saturation is not needed for the equivalence of divisibility: $H_k = e^{o(L)} < p$

and $p \nmid P$, so $p \mid \mathfrak{G} \iff p \mid G$ on full rank. Smith normal form therefore does not create extra independent constraints at the target prime. For general background on Smith normal form and random integer matrices and cokernels see [3, 1]; the conditional statement here was proved separately rather than deduced from a universality theorem with independent entries.

7.1. **Independent auxiliary columns.**

Corollary 7.2 (Saturation and resilience with independent noise)**.** *In the prime-block matrix model of Theorem 6.1, let $Z_1, \dots, Z_m$ be independent identically distributed random variables modulo q, independent of every matrix entry. Replace $\mathcal{E}_q$ by the event that $N_1+Z_1, \dots, N_m+Z_m$ are equal modulo q. The same conditional saturation bound holds. For $q = p \in (X, 2X]$ the same conditional rank-resilience bounds hold, with the stated choice of C depending on ε.*

Proof. Use the witness argument in Theorem 6.1. Conditional on all unselected columns and their noise, and on the two selected columns in the witness rows, the candidate annihilator is fixed. The necessary equality of the other $m-2$ noisy columns has probability $\sum_x \nu_q(x)^{m-2}$, where $\nu_q = \mu_q * \mathrm{Law}(Z_1)$. The remaining entries of the selected two columns retain their unconditional law; their noise need not be exposed or used in estimating annihilation. The same atom and Hoeffding bounds thus apply, and the ratio of the two moments of ν_q is at most q^2 by the same Jensen argument. This proves both assertions. In particular, at a good prime one can condition on equality of labels of the full block F, choosing Z_i to be the A-block reciprocal sum multiplied by P modulo p. Theorem 4.1 and this corollary then hold simultaneously in that conditional probability space (without imposing the extra coordinate exposures for the rank assertion). □

7.2. **Small moduli and large-prime tails.** We give precise bounds for the typical gcd, taking $3 \le m \le \lceil (UV)^2 \rceil + 2$. All implicit errors in the following statements are uniform in this range.

Lemma 7.3 (Small-modulus uniformity)**.** *There are fixed $a_0, c > 0$ such that, for $Y = \exp(a_0 L^{1/3})$, every integer $2 \le q \le Y$ with $(q, P) = 1$ satisfies*

$$\mu_q(x) = q^{-1}(1 + O(e^{-cL^{1/3}}))$$

uniformly in x. The assertion remains true after removing any set of at most $n/4$ denominators from B, taking q coprime to the remaining product and using that product in the numerator.

Proof. Multiplication by the unit P does not affect uniformity; consider $\sum C_b b^{-1}$ modulo q. For $1 \le h < q$, let u_b be the least absolute residue of hb^{-1}. If $|u_b| \le \eta q$, then $bu_b = h + jq$ with $|j| \le \eta q_p + 1$. All such primes b divide one of these $O(\eta q_p + 1)$ nonzero integers. Each has at most $O((\log q + \log q_p)/\log q_p)$ prime factors in B, since its size is $O(qq_p)$. Choose $\eta = c_0/(\log q + \log q_p)$ for small fixed c_0. The number of these primes is less than $n/4$ for all $q \le Y$, because $q_p/\log q_p \asymp n$ and $\log q \ll L^{1/3}$. Thus $E_B(q,h) := \sum_b \|hb^{-1}/q\|^2 \gg n/(\log q + \log q_p)^2$. The difference transform is bounded by $\exp(-4E_B(q,h))$. By making a_0 small, the sum of these bounds over $h \neq 0$ is at most $e^{-cL^{1/3}}$ uniformly. Fourier inversion proves the estimate. After removing at most $n/4$ denominators at least $n/2$ of the original primes still have phase greater than η. The same Fourier bound applies with uniform constants, provided q is prime to the remaining product; the count of small phases is made only among those remaining primes. □

For q as above, column independence gives $\mathbb{P}(q \mid G) = S_m(q) = q^{1-m}(1 + O(me^{-cL^{1/3}}))$. The saturation error $e^{-c'L}$ is relatively negligible here because $m \log Y = o(L)$. Thus on the full-rank event the same formula holds for $q \mid \mathfrak{G}$, including prime-power moduli and simultaneous divisibility conditions with product at most Y.

For $b \in B$, divisibility by b of G is equivalent to a constant row $(c_{b1}, \dots, c_{bm})$. Let

$$r_m = 2^{-m} + 2 \cdot 4^{-m}.$$

Then $\mathbb{P}(b \mid G) = r_m$, and the events for distinct $b \in B$ are independent. Moreover, for $b^{j-1} \le Y$,

$$\mathbb{P}(b^j \mid G) = r_m b^{-(j-1)(m-1)}(1 + O(me^{-cL^{1/3}})). \tag{11}$$

Indeed condition on the constant row, remove it, and divide the remaining differences by b; the divisibility condition is then equality of m independent reduced numerators modulo b^{j-1}, to which Lemma 7.3 applies. When $j = 1$ it is exact. The same relative formula for full-rank $\mathfrak{G}$ follows from saturation and $m \log Y = o(L)$.

For large primes we use the bound $\beta(q)$ in Lemma 5.1 with $B' = B$. Note that $\beta(q)$ is nonincreasing with q (its parameter t is nondecreasing). If $Y > H_k, q_p$, then for $q \ge Y$ coprime to P it applies to μ_q. On full rank the first two columns are distinct and $N_2 \ne N_1$. Any common prime factor $\ell \ge Y$ must divide $N_2 - N_1$; this nonzero integer has magnitude $< P$, and so at most $\log P / \log Y$ such prime divisors. Conditional on the first two columns, the other $m - 2$ independent columns must coincide with their value modulo ℓ, with probability at most $\beta(Y)^{m-2}$. For $Y = \exp(a_0 L^{1/3})$, Stirling gives $\beta(Y) \le e^{-cL^{1/3}}$. By Theorem 7.1, for those primes $\ell \mid \mathfrak{G}$ is equivalent to $\ell \mid G$. We conclude

$$\mathbb{P}(\operatorname{rank} M = k,\ \exists \ell \ge Y : \ell \text{ prime},\ \ell \mid \mathfrak{G}) \le \frac{\log P}{\log Y}\beta(Y)^{m-2} \le e^{-c'(m-2)L^{1/3}}. \tag{12}$$

Constants can be decreased to absorb the prefactor for the stated range of m.

Theorem 7.4 (Typical common divisor and Poisson limit)**.** *Define the random squarefree product $G_B = \prod_{b \in B:\, c_{b1} = \cdots = c_{bm}} b$. Then*

$$\mathbb{P}(\operatorname{rank} M = k,\ \mathfrak{G}(M) = G_B) = 1 - O(2^{1-m}).$$

For $m = \lceil (UV)^2 \rceil$, in particular,

$$\mathbb{P}(\operatorname{rank} M = k,\ \mathfrak{G}(M) > 1) \sim n2^{-m}.$$

If $m \to \infty$ in the stated range and $n2^{-m} \to \lambda \in (0, \infty)$ along a sequence, the number of distinct prime factors of $\mathfrak{G}$, defined as zero on rank failure, converges in law to $\mathrm{Pois}(\lambda)$.

Proof. Exclude integer lattice failure, of probability e^{-cL} by Theorem 6.1. By the exact identity of Theorem 7.1, it suffices to prove the assertion for G. Exclude all prime divisors not in B below Y, whose total probability by Lemma 7.3 is at most $(1 + o(1)) \sum_{j \ge 2} j^{1-m} = O(2^{1-m})$. Exclude squares of the denominator primes, whose probability by (11) is $O(nr_m(\min B)^{1-m}) = o(2^{1-m})$. Equation (12) excludes larger prime factors with probability $o(2^{1-m})$. Thus $G = G_B$ except for the stated error. Since G_B has independent prime-inclusion indicators of parameter r_m, its number of factors is exactly binomial (n, r_m). At $m = \lceil (UV)^2 \rceil$ we have $nr_m \to 0$ and $\mathbb{P}(G_B > 1) \sim n2^{-m}$; the error is negligible. In the second regime the usual binomial Poisson limit and the vanishing error prove the conclusion. □

The range condition on m is essential for the relative small-modulus formulas as used here; the proof is not a claim of universality for unrestricted growing matrix widths. The typical gcd theorem is also not a precise estimate for an individual very large prime dividing the gcd.

7.3. **Conditional gcd laws at moderately growing moduli.** The small-modulus estimate and the large-prime Hamming tail can be combined without dividing a weak absolute estimate by an uncontrolled conditioning probability. In particular, the deterministic imposed factor can be separated from the random denominator-prime product.

Theorem 7.5 (Conditional typical gcd and Poisson law)**.** *There is a constant $a_1 > 0$ such that, uniformly for $3 \le m \le \lceil (UV)^2 \rceil + 2$ and integers $1 \le q_0 \le \exp(a_1 L^{1/3})$ with $(q_0, P) = 1$, one has*

$$\mathbb{P}(\operatorname{rank} M = k,\ \mathfrak{G}(M) = q_0 G_B \mid \mathcal{E}_{q_0}) = 1 - O(2^{1-m}).$$

Here $\mathcal{E}_1$ means the sure event. Put $K_B = \#\{b \in B : c_{b1} = \cdots = c_{bm}\}$ and $\epsilon_L = e^{-c_2 L^{1/3}}$ for a sufficiently small constant $c_2 > 0$. For each integer $0 \le r \le n/4$,

$$\mathbb{E}\left[\binom{K_B}{r} \,\middle|\, \mathcal{E}_{q_0}\right] = \binom{n}{r} r_m^r (1 + O(m\epsilon_L)), \qquad r_m = 2^{-m} + 2 \cdot 4^{-m}.$$

Consequently, for $m \to \infty$ and $n2^{-m} \to \lambda \in (0, \infty)$, the number of distinct prime factors of $\mathfrak{G}/q_0$ on full rank (set to zero otherwise) converges conditionally in law to $\mathrm{Pois}(\lambda)$. *For $m = \lceil (UV)^2 \rceil$, uniformly in the imposed modulus,*

$$\mathbb{P}(\operatorname{rank} M = k,\ \mathfrak{G}(M) > q_0 \mid \mathcal{E}_{q_0}) \sim n2^{-m}.$$

Proof. For $q_0 = 1$ the statements about gcd follow from Theorem 7.4; the moment identity is exact. We give the conditional argument, which also accommodates this case with $S_m(1) = 1$. Let $Y = \exp(a_0 L^{1/3})$ from Lemma 7.3 and set $Y' = Y^{1/2}$. Choose $a_1 < a_0/2$ sufficiently small below. The conditioning probability is $S_m(q_0) = q_0^{1-m}(1 + O(m\epsilon_L))$. The saturation failure probability conditional on $\mathcal{E}_{q_0}$ is exponentially small in L by Theorem 6.1. On saturation we may use G instead of $\mathfrak{G}$.

If $\ell \notin B$ is a prime $\ell \le Y'$, the conditional event that $q_0\ell \mid G$ has probability

$$\frac{S_m(q_0\ell)}{S_m(q_0)} = \ell^{1-m}(1 + O(m\epsilon_L)).$$

This applies also to primes dividing q_0, detecting an extra power beyond the imposed exponent. It follows by small-modulus uniformity since $q_0\ell \le Y$. Summing over these primes gives $O(2^{1-m})$. On full rank $G > 0$, so any extra prime not in B and larger than Y' belongs to the large-prime tail. Replacing Y by Y' in the Hamming argument for (12), applied to G itself, bounds its unconditional probability on full rank by $e^{-c_3(m-2)L^{1/3}}$ for a fixed $c_3 > 0$. Choose $a_1 < c_3/4$; dividing by $S_m(q_0) \ge q_0^{1-m}$ gives at most $e^{-c_3(m-2)L^{1/3}/2} = o(2^{1-m})$ (since $(m-1)/(m-2) \le 2$).

For $b \in B$ the event $b^2 \mid G$ implies a constant b-row. Conditional on that row, differences of the original numerators equal b times differences of numerators formed with P/b and the remaining columns. The condition $q_0 \mid G$ together with $b^2 \mid G$ is therefore coincidence of the reduced columns modulo $q_0 b$, which is coprime to their remaining denominator product and at most Y for large X. Lemma 7.3 gives

$$\mathbb{P}(\{b^2 \mid G\} \mid \mathcal{E}_{q_0}) = r_m b^{1-m}(1 + O(m\epsilon_L)).$$

The sum over B is $o(2^{1-m})$, since $n \asymp L$ and $\min B \asymp LU$. Each first power $b \mid G$ is exactly the constant-row condition, with or without modular conditioning. The saturation error and the preceding bounds show that, except with conditional probability $O(2^{1-m})$, $G = q_0 G_B$ and the Smith identity identifies G with $\mathfrak{G}$.

For the moments fix r rows $E \subseteq B$. The event that they are all constant has unconditional probability r_m^r. Given their constant values the original numerator differences are the differences of the remaining reduced numerators multiplied by $\prod_{b \in E} b$, a unit modulo q_0. By the deletion-uniform form of Lemma 7.3, their coincidence probability is $q_0^{1-m}(1 + O(m\epsilon_L))$ uniformly, even for $r \le n/4$. Dividing by $S_m(q_0)$ and summing over the row sets proves the factorial moment formula. In the critical regime each fixed factorial moment converges to the corresponding Poisson moment, which determines the distribution; tightness follows from the first moment, and higher moments give uniform integrability. The conditional probability of any discrepancy between the two prime counts tends to zero by the gcd estimate. In the rare regime, $nr_m \to 0$, the first two moments give $\mathbb{P}(K_B > 0 \mid \mathcal{E}_{q_0}) \sim nr_m$ by the first two Bonferroni inequalities. The error $O(2^{1-m})$ is negligible compared to $nr_m \sim n2^{-m}$. This proves the last claim. □

The conditional theorem does not include target primes of size X: its imposed modulus is much smaller. It is a distribution theorem for a conditioned arithmetic matrix, not an estimate of the probability that a specified target prime divides its gcd.

8. Factorial moments of the number of common large primes

Define Ω_X to count primes $p \in (X, 2X]$ dividing $\mathfrak{G}(M)$ on full rank and set it to zero otherwise. Put $D_X = \lfloor \log P/L \rfloor = (\kappa + o(1))U$, $\kappa = C/10$, and for $r \geq 1$ let

$$Q_* = \lceil C \log(2X) \log_2(2X) \rceil, \qquad t_r = \min(n, \lfloor rL/\log Q_* \rfloor), \qquad s_r = \lfloor t_r/2 \rfloor,$$

$$\beta_r = \min\left(2^{-t_r}, \left[\sum_{j \leq s_r} \binom{n}{j}\right]^{-1}\right).$$

The common upper bound Q_* is independent of the target prime, and the proof of Lemma 5.1 works with this upper bound instead of q_p.

Proposition 8.1 (Factorial moment bound)**.** *For every $r \geq 1$,*

$$\mathbb{E}\binom{\Omega_X}{r} \leq \binom{D_X}{r} \beta_r^{m-2}.$$

In particular if $r = o(U)$ then $\log \beta_r^{-1} \geq (1/2 - o(1))(rL/U)\log(U/r)$. If $r = \lfloor \alpha U \rfloor$ with $0 < \alpha < \kappa$ fixed, then

$$\mathbb{E}\binom{\Omega_X}{r} \leq \exp\big(- [\kappa\, \mathsf{H}(\alpha/(2\kappa)) - o(1)] mL \big)$$

for $m \to \infty$, where $\mathsf{H}(x) = -x \log x - (1-x)\log(1-x)$.

Proof. On full rank $N_2 - N_1 \neq 0$, and all primes counted by Ω_X divide it. There are at most D_X candidates. Given the first two columns, a selection of r candidate primes has product $q \geq X^r$ coprime to P. Lemma 5.1 bounds the probability that the remaining columns all coincide with the first modulo q by β_r^{m-2}. Sum over candidate selections and then average over the first two columns. This proves the inequality (if $r > D_X$ it is zero). The two asymptotic statements follow from the binomial coefficient estimates $\binom{n}{s} \geq (n/s)^s$ and Stirling, respectively; $\log \binom{D_X}{r} = O(U)$ in the latter case. For sufficiently large fixed κ the entropy coefficient can exceed one. This controls aggregation of many large factors, not a single specified factor. □

For example, choose fixed $\alpha > 0$ and large C so that $\kappa \mathsf{H}(\alpha/(2\kappa)) > 1 + \delta'$ for fixed $\delta' > 0$. With $m = \lceil (UV)^2 \rceil$, since $\mathbb{P}(\mathcal{E}_p) \geq p^{1-m}$, Proposition 8.1 gives $\mathbb{P}(\Omega_X \geq \lfloor \alpha U \rfloor \mid \mathcal{E}_p) \leq e^{-cmL}$ for every target prime. This does not rule out the case $\Omega_X = 1$.

8.1. The relation to the high-order gcd average.

For clarity, the sufficient concentration target for pointwise coverage in [9] is

$$A_X := \max_{X < p \leq 2X} \log(p \rho_p(B)) = o(L/(UV)), \tag{13}$$

where the maximum is over primes and $\rho_p(B) = \max_x \mu_p(x)$ for the difference law. By nonnegative Fourier coefficients it equals $\mu_p(0)$. The high-order equivalence in [9] can now be obtained directly from Theorem 6.1. We recall it to make the limitation of the present gcd results precise.

Proposition 8.2 (High-order average, as in the companion expansion paper)**.** *For $m = \lceil (UV)^2 \rceil$,*

$$\frac{\log \mathbb{E}\Omega_X - (2-m)L}{m} = A_X + O(L/m + 1).$$

In particular, $\mathbb{E}\Omega_X \leq X^{2-m} \exp(o(mL/(UV)))$ is equivalent to (13).

Proof. By Theorems 6.1 and 7.1, conditional on $\mathcal{E}_p$ full rank holds with probability $1 - O(e^{-cL})$, and on full rank $p \mid \mathfrak{G}$ if and only if $p \mid G$. Thus

$$\mathbb{E}\Omega_X = \sum_{X < p \leq 2X} (1 + O(e^{-cL})) S_m(p).$$

The inequalities $\rho_p(B)^m \le S_m(p) \le \rho_p(B)^{m-1}$ follow from the definition. Retaining a prime attaining A_X gives the lower bound $\log \mathbb{E}\Omega_X \ge m(A_X - L - O(1)) + o(1)$, whereas summing at most $2X$ upper bounds gives $\log \mathbb{E}\Omega_X \le (2-m)L + (m-1)A_X + O(1)$. As $0 \le A_X \le L + O(1)$, the claimed identity follows. Finally $L/m = o(L/(UV))$. □

The unconditional prime-average bound $\sum_{X<p\le 2X} p\rho_p(B) \ll XU$ follows by counting the primes dividing nonzero reciprocal numerators, whose logarithmic size is $O(LU)$, with zero probability 2^{-n}. Thus a density-one conclusion is available, but the worst single prime in (13) is not controlled. The factorial moments in Proposition 8.1 constrain the number of different large divisors; they do not supply the exponentially precise first moment of the high-order collision event. The conditional gcd theorem above does not justify imposing a target-prime coincidence in the small-modulus formula.

9. A uniform distinct-value model: reduced heights and prime factors

The scale q now tends to infinity independently of X; all notation in this section is local. Put $\mathcal{H}(D) = R(D)$, $\mathscr{S}(D) = \{R(J) : J \subseteq D\}$, $\mathscr{S}_j = \mathscr{S}([1,j])$, $S(j) = |\mathscr{S}_j|$, $s(j) = \log S(j)$, $s(D) = \log|\mathscr{S}(D)|$ and $\mathcal{L}(j) = \operatorname{lcm}(1,\dots,j)$. For a real $P > 1$ let $b(P) = \min\{j : s(j) \ge \log P\}$, as in [7]. We need only the elementary large-prime decomposition and the four-point majorant, reproduced here for the height and mass estimates. Choose a sufficiently small absolute $c > 0$ and fix the coefficient bound $D \ge 4$ in the next lemma.

Put $\mathcal{L}(n) = \operatorname{lcm}(1,\dots,n)$. We use the prime number theorem, including the classical error term $\pi(x) = \operatorname{Li}(x) + O(xe^{-c\sqrt{\log x}})$, and $\log \mathcal{L}(n) \sim n$.

Lemma 9.1 (Bounded-coefficient rigidity)**.** *Fix an integer $D \ge 4$ and take a sufficiently small absolute $c > 0$. For all sufficiently large n let*

$$y_n = \frac{n}{c\log n}, \quad k_r = \lfloor n/r \rfloor, \quad V_r = r[1,k_r] \quad (y_n < r \le n), \quad K_n = [1,n] \setminus \bigcup_{r>y_n} V_r.$$

Here and below r in these expressions is prime. These blocks are disjoint. Any relation $\sum_{a\le n} c_a/a = 0$ with integer $|c_a| \le D$ splits into zero relations on each V_r and on K_n.

Proof. For large n, $y_n > \sqrt{n}$ and $k_r \le c\log n$. Choose c small enough that $D\mathcal{L}(k_r)H_{k_r} < r$ uniformly. Multiplication by r and reduction modulo r gives $\sum_{a\le k_r} c_{ra}/a \equiv 0 \pmod r$, since all other denominators are units at r. Multiplying by $\mathcal{L}(k_r)$ gives an integer divisible by r and of absolute value strictly less than r. It is zero. Subtract all these block relations from the original relation to obtain the core relation. □

The sum map from the product of local distinct sumsets is one-to-one and onto, by Lemma 9.1. Thus

$$s(n) = s(K_n) + \sum_{r>y_n} s(k_r). \tag{14}$$

The same assertion holds for intersections of all the components with a specified subset of $[1,n]$. It is essential here that the product is of *distinct rational values*, not of arbitrary representations of those values. The least common multiple of the core divides $\prod_{r\le y_n} r^{\lfloor \log n/\log r\rfloor}$. Its logarithm is $O_c(n/\log n)$, by the prime number theorem and the bound $O(\sqrt{n}\log n)$ for the contribution of higher powers. All core sums have this common denominator and size at most H_n, whence $s(K_n) = O_c(n/\log n)$.

Uniformly for $1 \le j \le \log n$, the prime number theorem with its error term gives

$$\pi(n/j) - \pi(n/(j+1)) = \frac{n}{j(j+1)\log n}\left(1 + O\left(\frac{\log_2 n}{\log n}\right)\right).$$

The relative error is uniform because the intervals have length at least a constant multiple of $n/(\log n)^2$. Using $s(j) \le j\log 2$ in (14), the normalized total relative-error contribution is $O((\log_2 n)^2/\log n) = o(1)$. Extending the upper summation limit from $c\log n$ to $\log n$ adds

only $O_c(1)$ to the normalized sum. A partially included terminal interval is similarly harmless. This proves the recursion $s(n) = \frac{n}{\log n}(\sum_{j\le \log n} s(j)/(j(j+1)) + O(1))$.

The trivial bound $s(j) \le j \log 2$ gives the first upper bound below; substituting it once more in the recursion gives the second. The lower bound follows from $s(j) \ge \pi(j)\log 2$, since sums of distinct primes' reciprocals are distinguished by reduction modulo each such prime. In detail, the sum of $1/(j\log j)$ up to $\log n$ is of order $\log_3 n$. Consequently

$$\frac{n\log_3 n}{\log n} \ll s(n) \ll \frac{n\log_2 n}{\log n}, \qquad s(n) \ll \frac{n}{\log n}(\log_3 n)^2. \tag{15}$$

For the second upper bound use $\sum_{2\le j\le \log n} \log\log j/((j+1)\log j) = O((\log_3 n)^2)$, handling small terms separately.

9.1. **The top-level entropy pool.** Write $\mathcal{G}_n = \bigcup_{r>y_n} V_r$. The exact recursion and (15) give

$$s(\mathcal{G}_n) = (1-o(1))s(n), \qquad \mathcal{H}(\mathcal{G}_n) = \sum_{r>y_n} \frac{H_{k_r}}{r} \ll \frac{(\log_2 n)^2}{\log n} = o(1). \tag{16}$$

For the mass estimate use $H_{k_r} = O(\log_2 n)$ and partial summation for primes in $(y_n, n]$.

Lemma 9.2 (Recursive two-element components)**.** *Apply the decomposition of Lemma 9.1 recursively to the children $[1, k_r]$, treating each core as a single component and stopping at a sufficiently large fixed height m_0. All sums factor into independent local distinct-value components after rescaling. For any collection of roots with sizes greater than m_0, the total number of two-element terminal components is at least a_0 times the total entropy, for an absolute $a_0 > 0$ depending only on the fixed choices.*

Proof. At every internal node of size $n > m_0$, the primes $n/3 < r \le n/2$ occur as markers and give at least $c_1 n/\log n$ two-element children (take m_0 sufficiently large). All its children number $O(n/\log n)$ and its core entropy is $O(n/\log n)$. Every terminal leaf entropy is bounded in terms of m_0, and every terminal leaf is the child of an internal node, since all roots exceed m_0. Summing over internal nodes, the total entropy is at most a constant times $\sum n/\log n$, whereas the number of two-element terminal leaves is at least a positive constant times the same sum. Exact factorization follows by applying Lemma 9.1 at every node. Rescaling denominators along the tree does not change the distinct-value counts. The same splitting applies to any zero relation assembled from component differences, as the coefficients in a difference of ordinary sums are bounded. □

For a large height q let

$$J_0 = \left\lfloor \sqrt{s(q)\log q/q} \right\rfloor, \qquad U = \bigcup_{y_q < r \le q/J_0} V_r.$$

Then $J_0 \to \infty$. The discarded top-level entropy is $O(J_0 q/\log q) = o(s(q))$ by grouping r according to $k_r < J_0$ and using $s(j) \le j\log 2$. Thus $s(U) = (1-o(1))s(q)$, and a common denominator D_U of its sums satisfies

$$D_U = \mathcal{L}(\lfloor c\log q\rfloor) \prod_{y_q < r \le q/J_0} r, \qquad \log D_U = O(q/J_0 + \log q) = o(q). \tag{17}$$

All retained children have sizes at least J_0; consequently Lemma 9.2 applies for large q. For any additional deletion of marker blocks of total entropy $o(s(q))$, the remaining roots still satisfy the same lemma.

Lemma 9.3 (Four-point majorant and global zero atoms)**.** *There exists a constant $\sigma_0 \in (1/2, 1)$ such that the uniform law μ_U of distinct sums of U has two probability majorants K_1, K_2 supported on $\mathscr{S}(U) - \mathscr{S}(U)$, satisfying on the real line*

$$\widehat{K_i} \ge 0, \quad |\widehat{\mu}_U| \le \sqrt{\widehat{K_1}\widehat{K_2}}, \qquad K_i(\{0\}) \le e^{-\sigma_0 s(q)} \quad (i = 1, 2). \tag{18}$$

The same conclusion holds after deleting from U a union of marker blocks of total entropy $o(s(q))$.

Proof. A two-element leaf has denominators $w, 2w$ for some integer w, so multiplication by $2w$ identifies its four values with $\{0,1,2,3\}$. The modulus of the Fourier transform of their uniform distribution at angle θ is $|\cos(\theta/2)\cos\theta|$. Define

$$F(\theta) = \frac{49 + 33\cos\theta + 15\cos 2\theta + 3\cos 3\theta}{100}.$$

This is the Fourier transform of a symmetric probability law on $[-3,3]\cap\mathbb{Z}$, with zero atom $49/100$. It majorizes the modulus for all real θ. To check this, let $z = |\cos(\theta/2)| \in [0,1]$. Then $F = (24z^6 - 6z^4 + 7)/25$. The finite certificate, as in the companion N02, can be checked without differentiation:

$$F - \frac{5}{18} = \frac{(6z^2-1)^2(12z^2+1)}{450},$$
$$\frac{2}{27} - z^2(1-2z^2)^2 = \frac{(2-3z^2)(6z^2-1)^2}{27}.$$

The first is nonnegative on $[0,1]$, and the second is nonnegative for $z^2 \le 1/2$. Since $2/27 < (5/18)^2$, this proves the required inequality on that range. For $z^2 \ge 1/2$ use

$$25(F - z(2z^2-1)) = (1-z)(7 - 2z^3 + 32(z - z^3) + 32(z^2 - z^3) + 24(z^3 - z^4) + 24(z^3 - z^5)),$$

whose right side is nonnegative on $[0,1]$. This proves the majorant, including nonnegativity.

For each two-element leaf take two copies of this law at its actual scale $1/(2w)$, both supported on the leaf difference set. For every other component with uniform law μ use the pair $(\mu * \widetilde{\mu}, \delta_0)$ or the reversed pair. Each pair dominates the modulus by the geometric mean of two nonnegative Fourier transforms. Convolve the first and second members respectively across the components. The zero atom of each global convolution is the product of the local zero atoms by rigidity of differences in Lemma 9.1. The negative logarithm of the product of the two zero atoms is at least

$$s(U) + 2a_0 s(U)\log(50/49) = (1 + c_2 - o(1))s(q), \quad c_2 > 0,$$

because an ordinary component contributes its entropy, whereas a two-element leaf contributes $-2\log(49/100) = \log 4 + 2\log(50/49)$. Allocate the ordinary components' nonzero capacities greedily to the smaller of the two accumulated capacities; the two-element leaves contribute equally to both. The remaining discrepancy is bounded by the maximum component entropy $O(\log q) = o(s(q))$, because every root of the forest has size at most $c\log q$. This yields (18) for a fixed $\sigma_0 > 1/2$ chosen sufficiently small. Deleting additional roots with entropy $o(s(q))$ does not change the argument. We fix such a $\sigma_0 < 1$ from now on. □

Here q tends to infinity independently of a target prime. Let x be uniform on $\mathscr{S}(\mathcal{G}_q)$ and write

$$H_q - x = u_x/v_x \quad \text{in lowest terms}, \qquad H_q - x = \mathcal{H}(K_q) + \sum_{r > y_q} \frac{b_r}{r}.$$

Complementation and exact factorization show that the b_r are independent uniform elements of $\mathscr{S}_{k_r}$. We use b_r for a local sum only, not for the inverse function $b(P)$. The rational number $H_q - x$ is positive and equals $\log q + O(1)$ uniformly by (16).

Theorem 9.4 (Height concentration)**.** *For each fixed $\varepsilon > 0$ there is $c_\varepsilon > 0$ such that*

$$\mathbb{P}(|\log v_x - \kappa_0 q| > \varepsilon q) + \mathbb{P}(|\log u_x - \kappa_0 q| > \varepsilon q) \ll e^{-c_\varepsilon q/\log q},$$

where $\kappa_0 = \sum_{j\ge 1}(1 - S(j)^{-1})/(j(j+1))$ and $5/8 < \kappa_0 < 3/4$.

Proof. For each top-level marker r, multiplication by r and reduction modulo r shows $r \mid v_x$ if and only if $b_r \neq 0$: a nonzero local value has numerator strictly between zero and r after multiplication by $\mathcal{L}(k_r)$, so cannot vanish modulo r. The exponent of r in v_x is then one. All other prime powers in v_x contribute at most $O(q/\log q)$ to its logarithm, by the core common-denominator estimate. Therefore

$$\log v_x = \sum_{r>y_q} I_r \log r + O(q/\log q), \quad I_r = \mathbf{1}_{\{b_r\neq 0\}}, \quad \mathbb{E}I_r = 1 - S(k_r)^{-1}.$$

For fixed j, the prime number theorem gives $\sum_{q/(j+1)<r\le q/j} \log r \sim q/(j(j+1))$. The total weight of $j > J$ is $O(q/J)$. First summing finitely and then letting J grow proves that the mean is $(\kappa_0 + o(1))q$. The sum of squares of the independent range lengths is $\sum_r (\log r)^2 = O(q \log q)$. Hoeffding's inequality then gives the stated fixed-error probability; the deterministic small-prime error and the asymptotic mean error are absorbed in the fixed tolerance. Since $\log u_x - \log v_x = O(\log_2 q)$, the same estimate holds for u_x. The upper bound on κ_0 follows from the first term and $1 - S(j)^{-1} < 1$; the lower bound uses $S(1) = 2$, $S(2) = 4$, $S(j) \ge 4$ for $j \ge 2$, and $S(3) > 4$. □

For a finite certificate of the numerical size of κ_0, let $Q_j = \mathcal{L}(j)$, $Q_0 = 1$, and $T_0 = \{0\}$. Define the integer sets recursively by

$$T_j = (Q_j/Q_{j-1})T_{j-1} \ \cup \ ((Q_j/Q_{j-1})T_{j-1} + Q_j/j).$$

Then $T_j = Q_j \mathscr{S}_j$, by splitting subsets according to membership of j, and $S(j) = |T_j|$. The exact integer counts through 22 are

j	1	2	3	4	5	6	7	8
$S(j)$	2	4	8	16	32	52	104	208
j	9	10	11	12	13	14	15	16
$S(j)$	416	832	1664	1856	3712	7424	9664	19328

j	17	18	19	20	21	22
$S(j)$	38656	59264	118528	126976	224128	448256

Let $C_J = \sum_{j\le J}(1 - S(j)^{-1})/(j(j+1))$. Monotonicity of S and telescoping give the exact interval

$$C_J + \frac{1 - 1/S(J)}{J+1} \le \kappa_0 < C_J + \frac{1}{J+1}.$$

In particular, using $J = 22$ and comparing rationals gives

$$\frac{69300497}{10^8} < \kappa_0 < \frac{69300508}{10^8}. \tag{19}$$

The computed tail-bracket width is $1/(23 \cdot 448256)$, less than 10^{-7}. The displayed rational endpoints have been rounded outward; no assertion of a numerical exact value is made.

The integer $\mathcal{L}(q)(H_q - x)$ has logarithmic height $(1 + o(1))q$ uniformly. It is not the same as the reduced numerator in Theorem 9.4.

For an integer $u > 0$, let $V_Y(u) = \sum_{\ell\le Y} v_\ell(u) \log \ell$, with the sum over primes and v_ℓ the usual valuation.

Theorem 9.5 (Small-prime mass)**.** *For the constant $\sigma_0 > 1/2$ in Lemma 9.3, and any fixed $0 < \sigma < \sigma_0$, uniformly for $Y \le e^{\sigma s(q)}$,*

$$\mathbb{E}V_Y(u_x) \ll q/J_0 + \log q + q\pi(Y)e^{-\sigma_0 s(q)} = o(q), \tag{20}$$

where $J_0 = \left\lfloor \sqrt{s(q)\log q/q} \right\rfloor$. Consequently the logarithmic mass of prime factors greater than $e^{\sigma s(q)}$ is $(\kappa_0 + o(1))q$ in probability.

Proof. Use the low-common-denominator pool U in (17), splitting $x = X_U + X_0$ into independent uniform components. Let K_i be its majorants in Lemma 9.3 and $C_i(\ell, a) = K_i\{z : v_\ell(z) \ge a\}$ with $v_\ell(0) = +\infty$. To use the majorization modulo a prime power, put $h = \lfloor \log q/\log \ell \rfloor$

and $Q = \mathcal{L}(q) = \ell^h Q'$ with $(\ell, Q') = 1$. For all candidate rational values, ℓ^h times the value is an ℓ-integer. The event $\ell^a \mid u_x$ is exactly $\ell^h(H_q - x) \equiv 0 \pmod{\ell^{a+h}}$ because the fraction is reduced and $a \geq 1$. A character of the finite cyclic group on $\ell^h z$ is a real-frequency character of z on this rational grid: multiplication of Qz by the inverse of Q' realizes the reduction. Thus the real-line majorization transfers to this group. Fourier inversion conditional on X_0, followed by Cauchy–Schwarz on the nonnegative majorants, gives

$$\mathbb{P}(\ell^a \mid u_x) \leq \sqrt{C_1(\ell, a) C_2(\ell, a)}.$$

The conditional translation only contributes a phase of modulus one. The sum of the Fourier coefficients of each majorant is the group order times its mass on zero, explaining the exact normalization.

Since $\log u_x = O(q)$, sum over $\ell \leq Y$, $1 \leq a \leq Cq/\log \ell$ for a sufficiently large C, and apply Cauchy–Schwarz with weights $\log \ell$. For each majorant the zero atom contributes at most $Cq\pi(Y)e^{-\sigma_0 s(q)}$. For any nonzero support value z of either majorant, the numerator of its reduced fraction has logarithmic size at most $\log D_U + O(1 + \log_2 q)$, as $D_U z \in \mathbb{Z}$ and $|z| \leq H_q$. The sum of all positive valuations weighted by $\log \ell$ is bounded by this logarithmic height. Therefore the nonzero part is $O(q/J_0 + \log q)$ by (17). This proves the first estimate in (20). Each term is $o(q)$ since $J_0 \to \infty$ and $\sigma < \sigma_0$. Markov's inequality and Theorem 9.4 prove the final assertion. □

For $q = \lfloor tb(P) \rfloor$ this controls the mass of prime factors up to $P^{\sigma t + o(1)}$. However, no lower bound for the proportion of these numerators divisible by a prime in the particular interval $[P, 2P]$ follows merely from their large-factor mass.

9.2. **A complementary high-order relation estimate.** For completeness we record the separate single-element-block argument, which uses a general concentration inequality in a cyclic group. It is not needed for Theorem 9.5, but applies to a broader class of finite linear sums. This Halász-type concentration estimate is a cyclic-group variant of the classical concentration method [5]; it uses Kneser's addition theorem [6].

Lemma 9.6 (A cyclic-group concentration estimate)**.** *Let $a_1, \ldots, a_n$ be units modulo $d \geq 2$, let $X = \sum_i \xi_i a_i$ with independent fair binary variables ξ_i, and let $R_k(d)$ be the number of ordered relations with k terms on each side, repetitions allowed. There exist absolute $C, c, \delta > 0$ such that, for $1 \leq k \leq \delta n$,*

$$\max_h \mathbb{P}(X = h) \leq d^{-1} + C\sqrt{k/n}\, R_k(d)/n^{2k} + e^{-cn}. \tag{21}$$

Proof. For $z \bmod d$ put $F(z) = \sum_i \|a_i z/d\|^2$, $T_u = \{z : F(z) \leq u\}$, and $A = R_k(d)/n^{2k}$. Fourier inversion bounds the probability by $d^{-1}\sum_z e^{-2F(z)}$. On T_u, the real part of $\sum_i e(a_i z/d)$ is at least $n - 2\pi^2 u$. Parseval for the $2k$-th power gives $|T_u|/d \leq Ae^{C_0 ku/n}$ for $0 \leq u \leq c_0 n$, with absolute constants and c_0 small. Every nontrivial subgroup G of the cyclic group satisfies $|G|^{-1}\sum_{z \in G} F(z) \geq c_1 n$ for an absolute $c_1 > 0$: multiplication by any unit permutes G, and a uniform grid of at least two points has mean squared distance from integers bounded below. Take $c_0 < c_1$. Then $T_{c_0 n}$ contains no nontrivial subgroup.

Let $U = c_0 n/k$. For $0 < u \leq U$ put $j = \left\lfloor \sqrt{U/u} \right\rfloor$. Minkowski's inequality gives the sumset inclusion $jT_u \subseteq T_U$. Since $0 \in T_u$, the stabilizer of jT_u must be trivial: a nontrivial stabilizer would itself lie in jT_u. Kneser's theorem yields $|jT_u| \geq j(|T_u| - 1) + 1$. Equivalently one may iterate Kneser, since every intermediate stabilizer is also contained in the final stabilizer. Applying the moment estimate at U gives $|T_u|/d \leq d^{-1} + C\sqrt{ku/n}\, A$. Finally use

$$d^{-1}\sum_z e^{-2F(z)} = \int_0^\infty 2e^{-2u}|T_u|/d\, du.$$

On $[0, U]$ use this last estimate, on $[U, c_0 n]$ use the moment estimate with $C_0 k/n \leq 1$, and on the rest use $|T_u| \leq d$. The middle integral is $O(e^{-U}A) = O(\sqrt{k/n}A)$, and the last is exponentially small in n. Choosing δ sufficiently small proves (21). □

Let n_1 count the primes $q/2 < r \le q$, so $n_1 \sim q/(2\log q)$. Their Bernoulli contributions $1/r$ are independent parts of the same uniform top-level candidate law. At a prime power ℓ^a omit the possible $r = \ell$ and condition on the rest of the non-Bernoulli components and that omitted variable. If their sum has negative ℓ-adic valuation, adding terms with unit denominators cannot cancel it, so the numerator divisibility event has probability zero. Otherwise apply Lemma 9.6 to inverse residues of the remaining r modulo ℓ^a. For $k \le \delta n_1/2$, the normalization by n_1 in place of $n_1 - 1$ changes only an absolute factor.

In the k-term relation count, a rational zero difference of reciprocal sums of primes implies equality of the multisets on each side: multiply by the distinct prime denominators and reduce at each, using $k < r$. There are at most $k!n_1^k$ such relations. A nonzero rational difference Δ becomes a nonzero integer after multiplication by the product of its distinct prime denominators; the logarithmic size of this integer is $O(k\log q)$. If ℓ is distinct from the denominators, its possible prime-power divisibilities have total logarithmic weight at most this bound. Summing (21) over ℓ^a, with weights $\log\ell$ and $a\log\ell \le Cq$, gives

$$\begin{aligned}\mathbb{E}V_Y(u_x) \ll\ & \log(2Y) + q\pi(Y)e^{-cn_1} \\ & + \sqrt{k/n_1}\left(q\pi(Y)\frac{k!}{n_1^k} + k\log q\right).\end{aligned}$$

Here the first term uses $\sum_{\ell\le Y}\log\ell/(\ell-1) \ll \log(2Y)$. The weighted sum for nonzero rational differences is taken after fixing the ordered tuple, so it requires no assumption of independence across moduli. In particular, for $\log Y = o(n_1)$, choose $k = \lfloor \eta n_1\rfloor$ with fixed small η, then let $\eta \downarrow 0$ after $q \to \infty$. The bound is $o(q)$, since $k!/n_1^k \le \eta^k$ and the last nonzero-difference term is $O(\eta^{3/2}q)$. Theorem 9.5 is stronger for the particular harmonic candidate law.

10. Fixed-width conditional zeta law

The estimate $O(2^{1-m})$ in Theorem 7.5 need not vanish when m is fixed. In this regime there is a nondegenerate residual gcd distribution. Keep the prime-block model and the same allowed moduli of that theorem, in particular $1 \le q_0 \le \exp(a_1L^{1/3})$, $(q_0, P) = 1$, with a_1 chosen there. Here $m \ge 3$ is fixed, $k = m - 1$. On the conditional space $\mathcal{E}_{q_0}$ define a positive integer R_X by

$$R_X = \frac{\mathfrak{G}(M)}{q_0G_B} \quad \text{if } \Lambda(M) = \mathbb{Z}^k, \qquad R_X = 1 \quad \text{otherwise.}$$

This is well-defined: on saturation $\mathfrak{G} = G > 0$, $q_0 \mid G$, $G_B \mid G$ by the constant-row characterization, and $(q_0, G_B) = 1$. Assigning a value on the exceptional event only avoids an undefined gcd ratio.

Corollary 10.1 (Uniform fixed-width residual gcd law)**.** *For every fixed $m \ge 3$ and every fixed positive integer d,*

$$\mathbb{P}(R_X = d \mid \mathcal{E}_{q_0}) \ \longrightarrow\ \frac{1}{\zeta(m-1)d^{m-1}}, \tag{22}$$

uniformly over the allowed imposed moduli q_0. In fact convergence of the conditional law to this probability law on the positive integers holds in total variation, uniformly over q_0.

Proof. For fixed d, eventually $(d, P) = 1$ and $q_0d \le Y$ uniformly, since the prime block moves to infinity and $a_1 < a_0/2$ in Theorem 7.5. On saturation, $d \mid R_X$ is equivalent to $q_0d \mid G$, also when d shares factors with q_0. The conditional probability of nonsaturation is $o(1)$ uniformly by Theorem 6.1. Since $\mathcal{E}_{q_0d} \subseteq \mathcal{E}_{q_0}$, Lemma 7.3 therefore gives the finite-divisibility limit

$$\mathbb{P}(d \mid R_X \mid \mathcal{E}_{q_0}) = \frac{S_m(q_0d)}{S_m(q_0)} + o(1) = d^{-k} + o(1). \tag{23}$$

As usual $S_m(1) = 1$. The errors are uniform over q_0 for each fixed d.

For an integer $A \geq 2$, the probability of a prime divisor of R_X larger than A satisfies

$$\limsup_{X\to\infty} \sup_{q_0} \mathbb{P}(\exists \ell > A : \ell \text{ prime},\ \ell \mid R_X \mid \mathcal{E}_{q_0}) \leq \sum_{j>A} j^{-k}. \tag{24}$$

Indeed on saturation an extra prime $\ell \notin B$ with $A < \ell \leq Y' = Y^{1/2}$ requires $q_0\ell \mid G$. By the uniform conditional ratio in the proof of Theorem 7.5 its probability is $\ell^{-k}(1 + O(m\epsilon_L))$, including extra powers of primes already dividing q_0. Sum this over the primes, bounding by the corresponding sum over integers. The large-prime tail for G on full rank beyond Y', divided by $S_m(q_0)$, is $o(1)$ uniformly by the same proof and choice of a_1. Finally a residual denominator prime $b \in B$ requires $b^2 \mid G$, because its first power has been removed and q_0 is coprime to P. The conditional square bound there yields total probability $O(nr_m(\min B)^{-k}) = o(1)$ for fixed $k \geq 2$. Off saturation $R_X = 1$ contributes no prime. This establishes (24) without using the nonvanishing fixed-width error $O(2^{1-m})$ as an $o(1)$ error.

Let P_A be the product of all primes at most A. For fixed d, A, the event $d \mid R_X$ but $d\ell \nmid R_X$ for every prime $\ell \leq A$ has, by finite inclusion-exclusion and (23), limiting conditional probability

$$\sum_{u \mid P_A} \mu_{\mathrm{Mob}}(u)(du)^{-k} = d^{-k} \prod_{\ell \leq A} (1 - \ell^{-k}),$$

where μ_{Mob} is the Möbius function. This event contains $R_X = d$; any excess event requires a prime $\ell > A$ dividing R_X/d, hence dividing R_X. The difference is bounded by (24). As $k > 1$, let $A \to \infty$ after $X \to \infty$, using the absolutely convergent Euler product for $\zeta(k)$, to obtain (22) uniformly. The candidate limit sums to one. To obtain uniform total variation convergence, choose D so that the sum of the limit probabilities for $d > D$ is arbitrarily small; finite uniform convergence for $d \leq D$ controls both the finite sum of absolute differences and the remaining mass uniformly. This proves the claim. □

The usual zeta law of gcds of independent uniformly sampled integers has the same limiting mass function; the new conditional feature here is the arithmetic separation of q_0 and the random product G_B. In particular the fixed-m residual probability of one tends to $1/\zeta(m-1)$, not to one. The argument gives no assertion at a target modulus of order X, and no claim of independence of the entire constant-row pattern from the residual gcd.

Scope and open questions

The two random models have different purposes. The conditional Gaussian theorem provides macroscopic peak statistics on a density-one set of primes, whereas the Smith identity and the rank bound hold even for the specified exceptional target primes. The new conditional gcd theorem applies to moderately growing moduli prime to the denominator product. The rare-event and Poisson conclusions there describe excess factors beyond the imposed modulus. The uniform distinct-value model concerns reduced positive numerators and denominators, not independent digit representations. A large logarithmic mass of large prime factors need not imply a hit in a particular dyadic interval. Thus neither the directed low-cost covering criterion of [7] nor the pointwise bound (13) of [9] is proved. The full exact local limit for the sparse conditional matrix and its possible arithmetic lattice corrections also remain separate questions. The results are internally reviewed mathematical derivations, not independently refereed or a claim of exhaustive priority verification.

References

[1] H. H. Nguyen and M. M. Wood, *Random integral matrices: universality of surjectivity and the cokernel*, Invent. Math. **228** (2022), 1–76. arXiv:1806.00596.

[2] J. E. Nymann, On the probability that k positive integers are relatively prime, *J. Number Theory* **4** (1972), 469–473, doi:10.1016/0022-314X(72)90038-8.

[3] R. P. Stanley, Smith normal form in combinatorics, *J. Combin. Theory Ser. A* **144** (2016), 476–495, doi:10.1016/j.jcta.2016.06.013.

[4] A. Ferber, K. Luh and G. McKinley, Resilience of the rank of random matrices, *Combin. Probab. Comput.* **30** (2021), 163–174, arXiv:1910.03619.
[5] G. Halász, Estimates for the concentration function of combinatorial number theory and probability, *Period. Math. Hungar.* (1977), doi:10.1007/BF02018403.
[6] M. DeVos, *A short proof of Kneser's addition theorem for abelian groups*, arXiv:1303.3539 (2013), Theorem 1.
[7] P. Chen, *Recursive Entropy and Optimal Modular Coverage of Harmonic Subset Sums*, unpublished manuscript, 2026.
[8] P. Chen and M. Ai, *Robust Multivariate Local Limit Theorems for Reciprocal Subset Sums*, unpublished manuscript, 2026.
[9] P. Chen, *Pointwise Expansion and Convolution Flattening of Modular Reciprocal Sumsets*, unpublished manuscript, 2026.