hexagon:2611.01297v1

POINTWISE EXPANSION AND CONVOLUTION FLATTENING OF MODULAR RECIPROCAL SUMSETS

PENGCHENG CHEN

Abstract. We study pointwise expansion and concentration of reciprocal subset sums in prime fields, with emphasis on the gap between rational entropy and modular entropy. A rank–height lifting theorem forces the denominators of rational points lying in a sufficiently small generalized arithmetic progression to divide a single small integer. A Boolean-cube sumset inequality yields an additive budget for progression covers of the actual reciprocal law; together with the Bogolyubov–Ruzsa and Balog–Szemerédi–Gowers theorems this gives logarithmic one-copy flattening gains of $(\log p)^{1/6-o(1)}$ for prime denominator blocks. Composite rectangular anchors give stronger heterogeneous expansion and, at every sufficiently large final prime, a near-full-capacity auxiliary convolution of order $O((\log_2 p)^{\varepsilon})$ for each fixed $\varepsilon > 0$. We establish independent packing and bounded-denominator obstructions to iterating these gains. We also give weighted packing inequalities for joint modular folding fibres in short windows. Using the recursive entropy formula for harmonic subset sums, we sharpen the growing reduced- and common-denominator entropy tradeoff to a uniform asymptotic equivalence. In particular, at height $Q \sim tb(p)$, where $s(b(p)) \sim \log p$ and $t > 1$, the least denominator budget p^K giving rational entropy $\log p$ satisfies $\log(Q/(K \log p)) = (\log_3 p)^{1-1/t+o(1)}$. These are pointwise rational-entropy and relative modular-expansion results, not a proof of uniform binary coverage at the counting scale. The remaining difficulty is the aggregate mass of all integer representatives of a residue when the common denominator is large.

Contents

1. Introduction and notation 2
2. Counting and spectral inputs 3
2.1. A spectral smoothing input at the larger height 4
2.2. A sufficient condition for bounded-cost coverage 4
3. Pointwise height rigidity and support expansion 5
3.1. Sharp entropy with varying bounded reduced denominators 6
3.2. A low-denominator block and exact mixed moments 9
3.3. A direct Poisson comparison and an exact CRT count 9
3.4. A rank–height rigidity theorem for modular reciprocals 10
3.5. Uniform support expansion by distinct denominator copies 12
4. Quasipolynomial covering obstruction and pointwise convolution flattening 13
4.1. Lifting rational points and bounding the mass of a union of progressions 13
4.2. A quasipolynomial cover from small doubling 15
4.3. A pointwise collision-probability ratio theorem 16
4.4. A quantitative relaxation of the covering criterion 16
5. Additive complexity budget from Boolean-cube sumsets 17
5.1. Improved pointwise expansion and robust convolution flattening 19
6. Hierarchical digit boxes and heterogeneous convolution 21
6.1. A rectangular block with its outer primes 21
6.2. An arbitrary-set mass inequality and asymmetric support growth 22
6.3. Symmetric flattening without a progression-cover theorem 23

Date: 28 November 2026.
2020 *Mathematics Subject Classification.* 11B75, 11D68, 11N25, 11T23, 60C05.
Key words and phrases. Reciprocal subset sums, generalized arithmetic progressions, convolution flattening, CRT, entropy, weighted packing.

6.4. A binomial-digit method for genuine asymmetric anticoncentration 24
6.5. A sharp digit-entropy obstruction at the smaller height 25
7. Density-one mixing on the auxiliary scale and the packing budget 26
7.1. A local lemma at denominator height proportional to the logarithm 26
7.2. Deterministic selection of the good auxiliary primes 28
7.3. An improved one-copy heterogeneous contraction 28
7.4. A budget obstruction to packing near-capacity blocks 29
7.5. Where a larger common denominator must enter 30
8. Growing denominator budgets and multiscale folding fibres 31
8.1. Rational entropy with a growing denominator budget 31
8.2. Joint folding in a window: a heterogeneous weighted packing bound 33
8.3. What the weighted-distance constraint cannot prove 34
9. An asymptotically sharp entropy budget at the recursive scale 35
9.1. Finite CRT counting interface for varying reduced denominators 36
9.2. Uniform asymptotic evaluation of the finite ledger 37
10. The pointwise covering problem and limitations 39
10.1. An equivalent high-order common-divisor question 39
10.2. Three limitations 40
10.3. Research direction 41
Scope of the results 41
References 41

1. Introduction and notation

For a finite set D of positive integers put $R(D)=\sum_{a\in D}1/a$, including $R(\varnothing)=0$. Let $\mathcal{V}(x)=\{R(J):J\subseteq[1,\lfloor x\rfloor]\cap\mathbb{Z}\}$, $H_q=\sum_{a\le q}1/a$, and $s(x)=\log|\mathcal{V}(x)|$. For a prime p and a denominator set $D\subseteq[1,p-1]$, write

$$\Sigma_p(D)=\{R(J)\bmod p:J\subseteq D\},\qquad \mu_{p,D}=\mathrm{Law}\left(\sum_{a\in D}\epsilon_a a^{-1}\right),\qquad \rho_p(D)=\|\mu_{p,D}\|_2^2,$$

where ϵ_a are independent Bernoulli variables of parameter $1/2$. Thus $\Sigma_p(D)$ is the support of $\mu_{p,D}$. We use counting-measure norms in physical space and the transform $\widehat{\mu}(h)=\sum_z\mu(z)\exp(2\pi ihz/p)$; Parseval has factor $1/p$. Put $\phi_{p,D}=\widehat{\mu}_{p,D}$. The independent differences $C_a=\epsilon_a-\epsilon'_a$ take $-1,0,1$ with probabilities $1/4,1/2,1/4$, and their transform is $|\phi_{p,D}|^2$. We use $|\cos(\pi v)|\le e^{-2\|v\|^2}$ and $\cos^2(\pi v)\le e^{-\sin^2(\pi v)}$. Write kS for a k-fold sumset; no repeated denominators are intended by this auxiliary operation. All modular fractions have unit denominators. A GAP will have formal volume equal to the product of its coordinate lengths, regardless of properness.

For a prime p, let $f(p)$ be the least $q<p$ for which $\Sigma_p([1,q])=\mathbb{F}_p$, and let $g(p)$ require in addition that every residue have a representative of real reciprocal mass at most one. Define $b(P)=\min\{n:s(n)\ge\log P\}$. The companion work [18] proves the entropy recursion, regular variation, and density-one optimal order of bounded-cost modular coverage. Its revised sufficient coverage factor $9200/4601$ is not used in the main results of the present paper: the symbol τ below generally denotes the growing energy scale $c_0L/(UV)$, not that fixed coverage factor. In the notation below these results give

$$b(p)\asymp\frac{\log p\log_2 p}{\mathcal{C}_4(p)},\quad f(p)\ge b(p),\qquad f(p)\le g(p)\ll\frac{\log p\log_2 p\log_3 p}{\mathcal{C}_5(p)}$$

for all sufficiently large primes, with $g(p)\ll b(p)$ on a density-one set of primes. The original height questions of Bleicher–Erdős and Erdős–Graham [1, 2], and subsequent work of Yokota and Liu–Sawhney [3, 4], motivate removing the residual factor uniformly. The question treated here is whether additional pointwise arithmetic information can accomplish that goal. This

paper proves genuine relative expansion estimates and sharp limitations on several strategies; it does not assume the unresolved absolute mixing estimate.

For fixed cutoff $A_* > 3$ let $\log_j$ denote the j-fold natural logarithm and set $\mathcal{C}_j(x) = \prod_{i\ge j:\,\log_i x\ge A_*} \log_i x$, with only defined iterates included. Empty products are one. We use the prime number theorem with its classical error term, $\log \operatorname{lcm}(1,\dots,n) \sim n$, and elementary Fourier and convolution inequalities. The following uniform comparisons will be used throughout:

$$\mathcal{C}_j(\log x) = \mathcal{C}_{j+1}(x), \qquad \mathcal{C}_j(y) \asymp \mathcal{C}_j(x) \quad (x^{1/2} \le y \le x^2,\ j \ge 2). \tag{1}$$

The first is an index shift. For the second, the second iterates differ by at most $\log 2$, and subsequent differences contract above the cutoff. The sum of logarithms of the ratios is bounded; at termination at most a bounded final factor is lost. In particular $\mathcal{C}_j(x) = (\log_j x)^{1+o(1)}$ for fixed $j \ge 2$. The constants are independent of the number of iterates. All fixed parameters are chosen before the prime tends to infinity.

The principal results are organized by mathematical mechanism. Theorem 3.4 and Theorem 4.2 give arithmetic obstructions to low-rank models; Theorems 5.1–5.3 give the strongest pure-prime relative gains. Theorems 6.2, 6.4, and 7.3 give composite-anchor estimates, and Theorems 7.5 and 7.7 quantify the budget restrictions. Theorems 8.1, 8.4, and 9.2 distinguish rational entropy from folding and sharpen the growing-budget theory. The final section records the uniform covering problem and precise sufficient conditions. All pointwise conclusions are uniform at their stated height scales, and no claim of independence of the actual inverse residues is used.

2. Counting and spectral inputs

The following entropy bound and dissociated-set estimate are included to specify the inputs from the counting-scale theory. The latter is the stronger set construction appearing in the proof of Bettin–Grenié–Molteni–Sanna, not a consequence merely of counting distinct rational sums.

Lemma 2.1. *For all sufficiently large real x,*

$$E(x) := \log|\mathcal{V}(\lfloor x\rfloor)| \ll \frac{x}{\log x}\mathcal{C}_3(x).$$

Proof. Partition denominators at prime factor threshold $y = x/\log x > \sqrt{x}$. The block with no prime factor exceeding y has a common denominator of logarithm at most $\vartheta(y)+O(\sqrt{x}\log x) \ll x/\log x$. Its sums lie on the corresponding rational grid in $[0, H_{\lfloor x\rfloor}]$. Every other denominator belongs to exactly one block $r[1,\lfloor x/r\rfloor]$, for $r > y$ prime. Consequently

$$E(x) \le C_0 x/\log x + \sum_{x/\log x < r \le x} E(x/r).$$

Write $u = \log x$. Partial summation with $\pi(z) = \operatorname{li}(z) + O(ze^{-c\sqrt{\log z}})$ gives

$$E(x) \le C_1 x/u + \frac{x}{u}(1 + O(\log u/u))\int_1^u \frac{E(v)}{v^2}\,dv.$$

For clarity, the prime sum is integrated against $d\pi(r)$ on $(x/u, x]$. The integrand $E(x/r)$ is monotone and bounded by $u\log 2$, so replacing the measure by $dr/\log r$ costs $O(xue^{-c'\sqrt{u}}) = o(x/u)$, including endpoints. Substitution $v = x/r$ and $\log r = u - \log v$ gives the displayed estimate. The error is uniform. The trivial bound $E(v) \le v\log 2$ immediately gives the elementary upper bound $E(x) \ll (x/\log x)\log_2 x$.

For the full upper bound, fix a large constant $v_0 > e^e$. Under the inductive hypothesis $E(v) \le K(v/\log v)\mathcal{C}_3(v)$ for $v \ge v_0$ below x, monotonicity of the product gives

$$\int_1^u \frac{E(v)}{v^2}\,dv \le C_{v_0} + K\mathcal{C}_3(u)(\log_2 u - \log_2 v_0)$$

for $u \geq v_0$. Here C_{v_0} is fixed. For sufficiently large x, $\mathcal{C}_3(x) = \log_2 u\, \mathcal{C}_3(u)$ exactly. The fixed positive subtraction $\log_2 v_0$ absorbs the $O((\log u/u)\log_2 u) = o(1)$ term and, for K large enough, the additive constants after division by $K\mathcal{C}_3(u)$. Thus the induction closes with one constant K, independently of the number of iterated logarithms. Formally apply induction over integer $\lfloor x \rfloor$ (all arguments $v \leq \log x$ are smaller); the majorant $(v/\log v)\mathcal{C}_3(v)$ is nondecreasing for large v, so the same bound holds at real arguments. Increase K on the finite initial range. This proves the lemma. □

Lemma 2.2 (Dissociated sets at the counting scale)**.** *For all large x there is a set $\mathcal{U}_x \subseteq [1,x]\cap\mathbb{Z}$ with pairwise distinct rational reciprocal subset sums, such that*

$$|\mathcal{U}_x| \gg \frac{x}{\log x}\mathcal{C}_3(x).$$

Proof. Use the set $\mathcal{U}$ defined in [5, Section 2]: $a \in \mathcal{U}$ if $1/a$ is not a signed sum with coefficients in $\{-1,0,1\}$ of reciprocals of smaller positive integers. Any nontrivial signed relation on $\mathcal{U} \cap [1,x]$ would contradict the defining property of its largest denominator. Thus the subset sums are distinct. The proof of [5, Theorem 1], specifically (7), (9), and (14) before applying Lemma 2, bounds $|\mathcal{U} \cap [1,N]|$ from below by

$$2\left(1 - \frac{1.4}{\log_k N}\right)\frac{N}{\log N}\prod_{j=3}^{k}\log_j N$$

when k is sufficiently large and $\log_k N \geq A_* > 3$. Take $N = \lfloor x \rfloor$ and the largest such k; the cutoff comparison (1) absorbs rounding. This uses the quantitative result of BGMS as an external input, not just the cardinality of $\mathcal{V}(x)$, which by itself would not ensure a dissociated subset of this size. □

2.1. **A spectral smoothing input at the larger height.** This paragraph isolates the only spectral input used in the conditional covering criteria below. It is the pointwise part of the rectangular-block estimates in [18, Uniform weak gap and pointwise rectangular estimate; source labels `lem:gap`, `prop:robust`]. For clarity, we give the precise formulation without relying on the numbering: for $X < p \leq 2X$, put $L = \log X$, $U = \log L$, $V = \log U$, $Q_0 = CLU$ and $q_p = \lfloor C \log p \log_2 p \rfloor$, where C is a sufficiently large fixed constant. Set $I_p = [q_p/2, q_p] \cap \mathbb{Z}$ and take

$$R_0 = \{r \text{ prime} : 8L < r \leq 12L\}, \qquad S = \{s \text{ prime} : CU/14 < s \leq CU/13\},$$
$$A = \{rs : r \in R_0, s \in S\}, \qquad T_S = \prod_{s\in S} s.$$

Then $A \subseteq I_p$ consists of distinct composites, $|A| \asymp_C L/V$, $R(A) = o(1)$, $\log T_S = O_C(U)$. Write

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

For every $h \neq 0$, $E_A(p,h) \gg_C L/(U^2V)$, and there is an absolute positive c_0 (with C fixed) such that the number of $h \neq 0$ with $E_A(p,h) \leq \tau_1 = c_0 L/(UV)$ is at most T_S. These two inputs follow by rational reconstruction and are established in N02. We use them only for the following conditional implication and its relaxations; none of the expansion, entropy-budget or folding theorems below assumes a near-uniform pointwise spectral bound. We fix these dyadic blocks in all conditional statements below.

2.2. **A sufficient condition for bounded-cost coverage.**

Proposition 2.3 (Support-set criterion)**.** *With τ_1 and A as above, suppose for a given $p \in (X, 2X]$ there is a set $D \subseteq I_p \setminus A$ such that*

$$|\Sigma_p(D)| \geq pe^{-4\tau_1+2}.$$

Then $g(p) \leq 3q_p$. In particular, a sufficient condition is $\rho_p(D) \leq p^{-1}e^{4\tau_1-2}$.

Proof. Let ν be uniform on $\Sigma_p(D)$ and μ the independent Bernoulli sum on A. By Parseval, $\sum_h |\widehat{\nu}(h)|^2 = p/|\Sigma_p(D)|$. Split the nonzero frequencies at $E_A(p,h) = \tau_1$. On the lower range, the two pointwise spectral estimates above bound the sum of squared convolution coefficients by $T_S e^{-cL/(U^2V)} = o(1)$. On the upper range it is at most $e^{-4\tau_1} p/|\Sigma_p(D)| \le e^{-2}$. Hence the squared coefficients including zero sum to less than 2. Cauchy–Schwarz shows that $S = \Sigma_p(A \cup D)$ has size greater than $p/2$. For every x, the sets S and $x - 3^{-1}S$ intersect, giving $S + 3^{-1}S = \mathbb{F}_p$. The blocks $A \cup D$ and $3(A \cup D)$ lie in disjoint intervals, and the total real reciprocal mass is at most $(4/3)H(I_p) = (4/3)\log 2 + o(1) < 1$, where $H(I_p) = \sum_{a \in I_p} 1/a$. All denominators are below p for large p. The collision-probability sufficient condition follows from $|\Sigma_p(D)| \ge 1/\rho_p(D)$. □

Lemma 2.4 (Weighted Hamming bound)**.** *Let B' be a set of distinct primes in I_p and write $n' = |B'|$. For any integer modulus $q \ge 2$ 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*

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

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 a prime block B in a fixed proportional subinterval of I_p, with $|B| \asymp L$, the lemma gives $\rho_p(B) \le \exp(-(1/2 - o(1))LV/U)$ at $q = p$, by the binomial lower bound and $t \sim L/U$. It applies without any unproved distribution assumption.

3. Pointwise height rigidity and support expansion

We first establish two pointwise results. The first determines the entropy obtainable from fractions of small *individual* reduced denominator, allowing that denominator to depend on the fraction. The second is a rank–height obstruction to generalized arithmetic progressions containing modular reciprocals. Combining the obstruction with the Green–Ruzsa theorem gives an unconditional support-expansion estimate. Neither result establishes the required nearly full entropy, and no exponential-scale inverse Littlewood–Offord theorem is assumed.

For this section use $L = \log p$, $U = \log_2 p$, $V = \log_3 p$, $W = \log_4 p$, and put

$$F_* = \mathcal{C}_4(p) = WG_*, \qquad G_* = \mathcal{C}_5(p).$$

The product identity holds for all sufficiently large p. In particular $F_* = V^{o(1)}$, $G_* = V^{o(1)}$. All comparisons below are uniform for integer Q in the range

$$c_* LU/F_* \le Q \le C_* LU, \tag{3}$$

where $0 < c_* < C_* < \infty$ are fixed. This includes both the counting scale and the larger scale used for the pointwise problem. Write

$$\mathcal{E}(Q,p) = \frac{QF_*}{UV}.$$

Then $\mathcal{E}(Q,p) \asymp L/V$ at the counting scale and $\mathcal{E}(Q,p) \asymp LF_*/V$ at the larger scale. Uniformly in (3), $\mathcal{E}(Q,p) = o(L)$.

3.1. **Sharp entropy with varying bounded reduced denominators.** For a rational number v, let $\mathrm{den}(v)$ denote its positive reduced denominator, with $\mathrm{den}(0) = 1$. For fixed $K > 0$ define

$$\mathcal{R}_K(Q,p) = \{v \in \mathcal{V}(Q) : \mathrm{den}(v) \le p^K\},$$
$$\mathcal{H}_K(Q,p) = \max_{D\subseteq[1,Q],\ \mathrm{lcm}(D)\le p^K} \log\big|\{R(J) : J \subseteq D\}\big|.$$

The empty least common multiple is one. In contrast to $\mathcal{H}_K$, the definition of $\mathcal{R}_K$ imposes no common denominator condition on the different rational values.

Theorem 3.1 (Variable-denominator entropy)**.** *For every fixed $K > 0$ and uniformly for (3),*

$$\log|\mathcal{R}_K(Q,p)| \asymp_{K,c_*,C_*} \mathcal{H}_K(Q,p) \asymp_{K,c_*,C_*} \mathcal{E}(Q,p). \tag{4}$$

Moreover, for any fixed $0 < \eta < 1$ there is a set $H \subseteq [\eta Q, Q] \cap \mathbb{Z}$ such that

$$|H| \asymp \mathcal{E}(Q,p), \qquad \mathrm{lcm}(H) \le p^{\min(K,1/4)}, \qquad R(H) = o(1), \tag{5}$$

whose reciprocal subset sums are distinct both as rational numbers and modulo p. All constants can depend on the fixed parameters and the cutoff.

Upper bound in (4). Let $y = L/V$. For large p one has $\sqrt{Q} < y < Q$. Put $D_Q = \mathrm{lcm}(1,\dots,Q)$ and factor it as

$$D_Q = T_0 \prod_{\substack{y<r\le Q\\ r \text{ prime}}} r,$$

where T_0 contains the full prime-power contributions from primes at most y. These factors are coprime, and

$$\log T_0 \le \vartheta(y) + O(\sqrt{Q}\log Q) = O(L/V).$$

Indeed, all additional prime powers have prime base at most $\sqrt{Q}$. For $v \in \mathcal{V}(Q)$ write $a_v = D_Q v \in \mathbb{Z}$. Since $r > y > \sqrt{Q}$, any denominator divisible by r is rm with $m \le Q/r < r$, and a denominator contains at most one such large prime. Therefore

$$a_v \pmod r \quad \text{has at most} \quad \exp(E(Q/r)) \quad \text{possible values},$$

where $E(x) = \log|\mathcal{V}(\lfloor x\rfloor)|$. This follows on reducing the numerator modulo r, since the local subset sum is a constant unit multiple of a reciprocal subset sum with denominators at most Q/r.

The key observation is that $r \mid \mathrm{den}(v)$ if and only if $a_v \not\equiv 0 \pmod r$. Hence for $v \in \mathcal{R}_K(Q,p)$ the support of the nonzero large-prime CRT coordinates has cardinality at most

$$s = \lfloor KL/\log y\rfloor = O_K(L/U).$$

The number of possible supports is at most $\sum_{j\le s}\binom{\pi(Q)}{j}$. Its logarithm is $O_K(LV/U)$: use $\pi(Q) = O(L)$ and the usual binomial bound (or $2^{\pi(Q)}$ if the number of available coordinates is at most $2s$). For each fixed support S, CRT permits at most

$$T_0 \exp\left(\sum_{r\in S} E(Q/r)\right)$$

residue patterns for a_v modulo D_Q. Since $0 \le v \le H_Q$, each such residue pattern gives at most $\lfloor H_Q\rfloor + 2$ distinct rational values. Thus

$$\log|\mathcal{R}_K(Q,p)| \ll_K \frac{L}{V} + \frac{LV}{U} + \log(H_Q+2) + \sum_{y<r\le L} E(Q/r) + sE(Q/L), \tag{6}$$

with the sum over primes. Here monotonicity of E bounds the contribution of the selected primes $r > L$; the primes between y and L have been included even when their coordinates vanish.

By Lemma 2.1 and the uniform product comparison (1), for $y < r \le L$ one has

$$E(Q/r) \ll (Q/r)G_*/V.$$

Indeed, Q/r lies between fixed multiples of U/F_* and UV, its logarithm is asymptotic to V, and $\mathcal{C}_3(Q/r) \asymp \mathcal{C}_3(U) = G_*$. Likewise $E(Q/L) \ll (Q/L)G_*/V$. The prime number theorem and partial summation give

$$\sum_{L/V<r\le L} \frac{1}{r} \ll \frac{\log V}{U} = \frac{W}{U}.$$

For example, replacing $d\pi(r)$ by $dr/\log r$ has an exponentially small relative-to-logarithms error on this range, and the resulting integral is $\log(U/(U-\log V))$. Consequently the last two terms of (6) are $O_K(QWG_*/(UV))$ and $O_K(QG_*/(UV))$, respectively. The other terms are $O_K(L/V)$, since $V^2 = o(U)$, and $L/V \ll \mathcal{E}(Q,p)$ in (3). This proves $\log|\mathcal{R}_K(Q,p)| \ll_K \mathcal{E}(Q,p)$. In particular, this argument does not assume that every prime appearing in a representing subset survives in the reduced denominator: primes that cancel simply have zero CRT coordinate. □

Construction for the lower bound. We use the dissociated-set estimate in the proof of Bettin–Grenié–Molteni–Sanna [5, Section 2, especially (7), (9), (14)]. For each sufficiently large M, there is a set $D(M) \subseteq [1,M]\cap\mathbb{Z}$ of size

$$|D(M)| \gg (M/\log M)\mathcal{C}_3(M) \tag{7}$$

with no nonzero signed reciprocal relation with coefficients in $\{-1,0,1\}$. To recall why the source provides the stronger dissociated-set statement rather than just the number of rational values: its set $\mathcal{U}(M)$ consists of integers $a \le M$ for which $1/a$ is not a signed sum of reciprocals of smaller positive integers. Taking the largest denominator in a hypothetical signed relation proves dissociation. The lower estimate on $|\mathcal{U}(M)|$ is obtained before the final use of the subset-value counting inequality. Choosing the last permitted logarithm at the fixed cutoff gives (7) with a uniform constant.

Using Lemma 2.1, we may also require $D(M) \subseteq [c_1M, M]$ for some sufficiently small fixed $c_1 > 0$, with the same order of cardinality. In fact, by dissociation, $|D(M)\cap[1,c_1M]|\log 2 \le E(c_1M)$, and monotonicity of the cutoff product and the upper estimate make this at most half of the lower bound in (7) for small enough c_1 and large M. We have the corresponding upper bound $|D(M)| \ll (M/\log M)\mathcal{C}_3(M)$.

Choose a fixed $\epsilon > 0$ with $(1+\epsilon)^{-2} > \eta$. Choose a large fixed ratio $A_0 > 4(1+\epsilon)^2/c_1$, and a sufficiently small fixed $a_0 > 0$ (depending on K, c_1). Set

$$Y_j = a_0 L A_0^{-j}, \quad M_j = \lfloor Q/Y_j \rfloor, \quad 0 \le j \le J = \left\lfloor \frac{\log V}{2\log A_0} \right\rfloor.$$

Uniformly in j, M_j lies between fixed multiples of U/F_* and $U\sqrt{V}$, so

$$\log M_j \sim V, \quad \mathcal{C}_3(M_j) \asymp G_*, \quad |D(M_j)| \asymp M_jG_*/V = o(U).$$

The last assertion follows from $G_* = V^{o(1)}$ and $M_j = O(U\sqrt{V})$. Partition $[c_1M_j, M_j]$ into a bounded number of intervals of fixed ratio at most $1+\epsilon$. Choose a subblock D_j of $D(M_j)$ in some $[s_j, (1+\epsilon)s_j]$ with

$$|D_j| \gg M_jG_*/V, \qquad c_1M_j \le s_j \le M_j.$$

Consider outer primes in the interval

$$\mathcal{P}_j = \left\{ r \text{ prime} : \frac{Q}{(1+\epsilon)^2 s_j} \le r \le \frac{Q}{(1+\epsilon)s_j} \right\}. \tag{8}$$

This interval has endpoints comparable to Y_j and fixed ratio $1+\epsilon$. All these intervals are pairwise disjoint by the choice of A_0: their lower endpoint is at least $Y_j/(1+\epsilon)^2$ and their upper endpoint is at most $2Y_j/c_1$. Moreover, $r \gg L/\sqrt{V} \gg \max_j M_j$. By the PNT, $|\mathcal{P}_j| \asymp Y_j/U$.

For each j discard those outer primes r at which some nonzero signed reciprocal relation on D_j vanishes modulo r. Each such rational relation is nonzero; after multiplication by $\mathrm{lcm}(1,\dots,M_j)$ its nonzero numerator has logarithmic absolute value $O(M_j)$ (use the harmonic

bound). It has at most $O(1+M_j/U)$ prime factors in $\mathcal{P}_j$. There are at most $3^{|D_j|}$ relations, hence the number of discarded primes is

$$3^{|D_j|}O(1+M_j/U)=\exp(o(U))=o(Y_j/U),$$

uniformly. Denote the remaining primes by $\mathcal{P}_j^*$; their local reciprocal subset sums on D_j are injective modulo r.

Set

$$H=\bigcup_{j=0}^{J}\{rd : r\in\mathcal{P}_j^*,\ d\in D_j\}.$$

The products lie in $[\eta Q,Q]$ by (8). They are distinct: an outer prime is larger than every inner denominator, and the outer-prime intervals are disjoint. Each scale contributes $\asymp (Y_j/U)(M_jG_*/V)\asymp QG_*/(UV)$ elements. Since $J+1\asymp W$,

$$|H|\asymp QWG_*/(UV)=\mathcal{E}(Q,p),\qquad R(H)\ll |H|/Q=O(F_*/(UV))=o(1).$$

A common denominator of H is

$$T=\operatorname{lcm}(1,\dots,\max_j M_j)\prod_j\prod_{r\in\mathcal{P}_j^*} r.$$

The logarithm of the first factor is $O(U\sqrt{V})=o(L)$. All the outer primes are at most $2a_0L/c_1$, so their product has logarithm at most $(2a_0/c_1+o(1))L$. With a_0 small enough, this proves $T\le p^{\min(K,1/4)}$.

Suppose two subset sums on H agree modulo p. Their rational difference has absolute value at most $R(H)<1$ and denominator dividing $T<p$. Thus its integer numerator after multiplication by T is divisible by p and has absolute value less than p; it is zero. Reducing the resulting equality modulo any outer prime r and dividing by the common factor prime to r gives a signed zero relation on D_j modulo r. The choice of good outer primes forces every coefficient in that block to be zero. Hence the two subsets are identical. The same argument gives rational dissociation. We have constructed $2^{|H|}$ different values all with denominator at most p^K, all arising from one set with common denominator at most p^K. This proves the lower bounds and (5). □

Corollary 3.2 (Pointwise support at both scales)**.** *For every fixed $0<\eta<1$, let $I_\eta=[\eta Q,Q]\cap\mathbb{Z}$ and let $\rho_p(I_\eta)$ be the collision probability of the independent Bernoulli reciprocal subset sum on this interval. Uniformly for all sufficiently large primes and* (3)*,*

$$\log|\Sigma_p(I_\eta)|\gg\mathcal{E}(Q,p),\qquad \rho_p(I_\eta)\le\exp(-c\mathcal{E}(Q,p)).$$

In particular, at $Q\asymp LU/F_*$ these are $\log|\Sigma_p(I_\eta)|\gg L/V$ and $\rho_p(I_\eta)\le\exp(-cL/V)$. At $Q\asymp LU$, the exponent improves to $\Omega(LF_*/V)$, with the full variable-depth product F_*.

Proof. Take H from Theorem 3.1. Its subset sum is uniform on $2^{|H|}$ residues; thus its maximum atom and collision probability equal $2^{-|H|}$. Independent convolution with the remaining denominators cannot increase either quantity, by convexity or Parseval. Its support is contained in the full support, which proves the other inequality. □

These estimates do not show that the entire reciprocal sumset has entropy $o(L)$: they concern only the small reduced denominators, or a particular injective block. But they exclude any plan to obtain p^c distinct values, for fixed $c>0$, by retaining only fractions of reduced denominator at most p^K (fixed K), even if the common denominator is allowed to vary with the represented value. This strictly strengthens a restriction on a single common denominator.

3.2. **A low-denominator block and exact mixed moments.** For later comparison with long folding fibres we retain the mixed-moment calculation from the pointwise theory. Work with the dyadic blocks of Section 2.1, with B the primes in $[4Q_0/5, 9Q_0/10]$, $n = |B| \asymp L$, and $P = \prod_{b\in B} b$. The construction in the proof of Theorem 3.1, using $L = \log X$ and $Q = \lfloor Q_0 \rfloor$ instead of the individual prime parameters, with any fixed $\eta > 1/2$ and with the outer-prime upper cutoff chosen below $L/6$, supplies a single $H \subset I_p$ for every prime in the dyadic block, with

$$d = |H| \asymp L\mathcal{C}_4(X)/V, \qquad T = \operatorname{lcm}(H) < X^{1/4}, \qquad R(H) \ll \mathcal{C}_4(X)/(UV) = o(1/U).$$

The outer primes used by H are $\ll L$ and exceed all its inner denominators; they can be chosen smaller than $L/6$. The common denominator T is coprime to P, and H is disjoint from A and B because of the different largest prime factors. Its rational subset sums are distinct modulo p. This improves the size of the earlier fixed-depth construction without changing its mixed-moment argument.

Proposition 3.3 (Dissociated block and mixed moments)**.** *Write* $K_B(h) = \sum_{b\in B} e^{2\pi i h b^{-1}/p}$ *and* $k_0 = \lfloor L/(8\log q_p)\rfloor$*. For each integer* $0 \le k \le k_0$*,*

$$\sum_{h \bmod p} |\phi_{p,H}(h)|^2 |K_B(h)|^{2k} \le p 2^{-d} k! n^k. \tag{9}$$

Consequently $\rho_p(I_p) \le \exp(-c_C L\mathcal{C}_4(p)/V)$ *at every sufficiently large prime.*

Proof. Expand (9) by orthogonality. Its left side is $p4^{-d}$ times the number of equalities modulo p between a subset of H plus an ordered k-tuple from B and another such choice. The denominator of their rational difference divides an integer at most $Tq_p^{2k} < p^{1/2}$; its real magnitude is $O(1)$, since $k = O(L/U)$ and the primes in B are of order LU. A modular equality is therefore a rational equality. Reducing at each prime of B gives equality of the occurrence counts ($k < b$ and $b \nmid T$); the two subsets of H must then coincide. The number of choices is at most $2^d n^k k!$. This proves the displayed bound. Split at $|K_B(h)| = n/2$. On the large range the normalized sum of squared H coefficients is bounded by $2^{-d}4^{k_0}k_0!/n^{k_0}$. On the small range the inequality $\Re K_B(h) \ge n - 2\pi^2 E_B(p,h)$ gives $E_B(p,h) > n/(4\pi^2)$, so $|\phi_{p,B}(h)|^2 \le e^{-n/\pi^2}$. Parseval and the disjointness of H and B now give

$$\rho_p(I_p) \le 2^{-d} e^{-cLV/U} + e^{-n/\pi^2}.$$

The claimed estimate follows. Alternatively the modular injectivity of H alone gives the same order of the exponent; the mixed-moment inequality additionally records exact low-order information. □

3.3. **A direct Poisson comparison and an exact CRT count.** For H from Proposition 3.3, let $u = T\sum_{a\in H} C_a/a$, with integer law $\omega(u)$. Independently for each $b \in B$ take $C_b^* = N_b^+ - N_b^-$ with $N_b^\pm \sim \mathrm{Pois}(1/4)$. Write $w(c) = \mathbb{P}(C_b^* = c)$. The Fourier comparison $\cos^2(\pi x) \le e^{-\sin^2(\pi x)}$ gives

$$\rho_p(H \cup B) \le \mathbb{P}\left(u/T + \sum_{b\in B} C_b^*/b \equiv 0 \pmod p \right).$$

Let $D_0 = TP$. For $z \in \mathbb{Z}$ set $u(z) = \langle pzP^{-1}\rangle_T$ and $c_b(z) = \langle pz(D_0/b)^{-1}\rangle_b$ (centred residues; for $T = 1$ use 0). With K large enough, define

$$\mathcal{W}_p = \sum_{\substack{|z|\le KD_0/(pU) \\ \sum_b |c_b(z)| \le n}} \omega(u(z)) \prod_b w(c_b(z)).$$

Then

$$\rho_p(H \cup B) \le \mathcal{W}_p + O(p^{-3}). \tag{10}$$

Indeed the Poisson tail $\sum |C_b^*| > n$ has probability at most $\exp(-cn) = O(p^{-3})$ for large C. On the truncated set $|u| < T/2$ and $|C_b^*| < b/2$ for large X. A collision is an integer equality

$Pu + \sum_b (D_0/b)C_b^* = pz$ with $|z| \le KD_0/(pU)$, and CRT uniquely recovers all coordinates. Conversely the same congruences imply this equality, since its error is a multiple of D_0 of absolute value $O(D_0/U) < D_0$. Thus $\mathcal{W}_p$ is exactly the truncated collision probability, not a heuristic independence model. Formula (10) is a valid alternate sufficient path, but has not been bounded by $p^{-1}e^{o(\tau_1)}$.

Two limitations of this route are useful. First, direct Fourier completion of the periodicised weight does not separate it from the original frequency problem. On $\mathbb{Z}/D_0\mathbb{Z}$, if its total probability is normalized to one, its Fourier coefficient at k is

$$\prod_{a\in H} \cos^2(\pi k\bar{p}/a) \exp\left(-\sum_{b\in B} \sin^2(\pi k\bar{p}/b)\right),$$

where $\bar{p} = p^{-1} \pmod{D_0}$. For $|k| < p$, the reciprocity identity $\bar{p}/b \equiv -b^{-1}/p + 1/(pb) \pmod 1$ converts the B factor to its finite-field analogue up to multiplicative $1 + O(1/U)$. The identity here uses b^{-1} modulo p. Second, for real CRT twists $\alpha_b = p^{-1} \pmod b$, the circle law of $\sum C_b\alpha_b/b$ has Fourier coefficient at $k = p$ equal to $\prod_b \cos^2(\pi/b) = 1 - o(1)$. The natural target is equality of peak masses, not a globally uniform circle law. Treating the twists as independent random residues would miss this mandatory resonance.

3.4. **A rank–height rigidity theorem for modular reciprocals.** For a prime $p > Q$, consider an affine generalized arithmetic progression (GAP) in $\mathbb{F}_p$,

$$\mathcal{P} = x_0 + \left\{\sum_{i=1}^{d} t_i x_i : 0 \le t_i < N_i,\ t_i \in \mathbb{Z}\right\}, \qquad \mathcal{L} = \prod_{i=1}^{d} N_i, \tag{11}$$

where N_i are positive integers. Here $\mathcal{L}$ is the *formal volume*; properness is not assumed. The integer heights of the coordinate vectors, not just the cardinality of the image, are essential. Finite integer certificate. The next module isolates the only determinant argument. It applies to rational inputs of bounded height as well as reciprocals. All its parameters are finite integers; no prime asymptotics, properness of the GAP, or inverse theorem enter it.

Finite module 1 (Integer certificate for lifting a rational column)**.** *Let C be an integer matrix with a nonempty finite row set I, with columns numbered $0, \dots, d$, whose zeroth column is all ones and with $|C_{ij}| \le N_j$ for $j \ge 1$, where $N_j \ge 1$. Put $\mathcal{L} = \prod_{j=1}^{d} N_j$, and let $r = \operatorname{rank}_{\mathbb{Q}} C \ge 1$. Write $v_i = a_i/b_i$ in lowest terms, with $a_i \in \mathbb{Z}$, $1 \le b_i \le H_b$, $|a_i| \le H_a$, where $H_a, H_b \ge 1$ are integers. Suppose p is prime, $p \nmid b_i$ for all i, and the reduced column $v \bmod p$ belongs to the span of $C \bmod p$. If*

$$(r+1)!H_b^r H_a \mathcal{L} < p, \tag{12}$$

then there is a positive integer T_0 such that

$$b_i \mid T_0\ (i \in I), \qquad T_0 \le r!H_b^r \mathcal{L} < p.$$

More precisely, take r rational basis columns J and r ordered basis rows I_0, write $C_* = C_{I_0,J}$, $D = \det C_* \ne 0$, and $B = \prod_{i\in I_0} b_i$. With $T = DB$ and $z = \operatorname{adj}(C_*)(Bv_{I_0}) \in \mathbb{Z}^r$, the following integer identities hold for all i:

$$b_i C_{i,J} z = Ta_i, \qquad T_0 = |T|. \tag{13}$$

Proof. Every r-minor of C has absolute value at most $r!\mathcal{L}$ by the Leibniz formula: each column contributes its bound at most once, including bound one for column zero. In particular $0 < |D| < p$. All $(r+1)$-minors vanish as integers because the rational rank is r, so the reduced matrix has rank r as well; if such minors do not exist, the rank upper bound is automatic. Thus the selected columns span the reduced column space. The vector Bv_{I_0} is integral. The adjugate identity gives $C_* z = DBv_{I_0} = Tv_{I_0}$, proving (13) on the basis rows.

For a row $i \notin I_0$, append it below the basis rows. Define the integer matrix M_i by taking $[C_J\ v]$ on these $r+1$ rows and multiplying each row by its own denominator b_h. Its last column

is then a_h. The determinant identity for a bordered matrix (or expansion in the last row) gives

$$\det M_i = Ta_i - b_i C_{i,J} z. \tag{14}$$

Indeed, before scaling the rows the bordered determinant is $Dv_i - C_{i,J}\operatorname{adj}(C_*)v_{I_0}$, and the row-scaling factor is $b_i B$. The modular span condition gives $\det M_i = 0 \pmod p$. Each term in its integer Leibniz expansion has exactly r factors bounded by the row-denominator bound H_b times their column bounds and one factor bounded by H_a. Therefore

$$|\det M_i| \le (r+1)! H_b^r H_a \mathcal{L} < p.$$

An integer divisible by p in this open interval must vanish. This proves (13) for every row. Coprimality of a_i, b_i implies $b_i \mid T$; the bound for $|T|$ follows from $|D| \le r!\mathcal{L}$ and $B \le H_b^r$. The sign of D is immaterial because the positive common denominator is $|T|$. □

Theorem 3.4 (Rank–height lifting)**.** *Let $D \subseteq [1, Q] \cap \mathbb{Z}$ be a nonempty set such that $a^{-1} \in \mathcal{P}$ for every $a \in D$. Suppose*

$$(d+2)!\, Q^{d+1}\mathcal{L} < p. \tag{15}$$

Then for some integer $1 \le r \le d+1$,

$$\operatorname{lcm}(D) \le r!\, Q^r \mathcal{L} < p. \tag{16}$$

In particular, if B' is a set of distinct prime denominators in $[b_-, Q]$, with $b_- > 1$, then the number of their modular inverses in $\mathcal{P}$ is at most

$$\frac{\log \mathcal{L} + (d+1)\log Q + \log((d+1)!)}{\log b_-} < \frac{\log p}{\log b_-}. \tag{17}$$

Proof. For each $a \in D$ choose one coordinate vector $(t_{a1}, \dots, t_{ad})$ giving a^{-1} in (11). Let the row of C be $(1, t_{a1}, \dots, t_{ad})$, and put $v_a = 1/a$. The GAP representation is exactly the modular column-span hypothesis of Finite module 1. Since $r \le d+1$, (15) implies (12) with $H_a = 1$, $H_b = Q$. The resulting integer T_0 is divisible by every $a \in D$, and $T_0 \le r!Q^r\mathcal{L}$. Thus $\operatorname{lcm}(D) \mid T_0$, proving (16). For distinct prime denominators the lcm is their product, which gives (17) by taking logarithms (as usual $b_- > 1$ in that assertion). If none of the prime inverses belongs to the GAP the cardinality bound is trivial. No properness or linear independence of the GAP generators is used. □

Corollary 3.5 (An explicit obstruction to inverse concentration models)**.** *Set $\tau = c_0 L/(UV)$ for any fixed $c_0 > 0$. For primes $b \asymp Q$ in a fixed proportional interval, under (3), a GAP of rank $d = o(\tau/U)$ and formal volume at most $p\exp(-a\tau)$, with any fixed $a > 0$, contains only $O(L/U)$ of their inverses. In particular it cannot contain a positive fixed proportion of a prime block of cardinality $\asymp Q/U$. More generally, if $D_0 \subseteq [1, Q]$ is rationally dissociated, then any GAP satisfying (15) contains the inverses of at most $O(\mathcal{E}(Q, p))$ elements of D_0. At the counting scale this is $O(L/V)$, so it cannot capture a positive proportion of a rationally dissociated set of cardinality $\asymp L$ there.*

Proof. The rank assumption gives $(d+1)\log Q + \log((d+2)!) = o(\tau)$, hence (15). The prime assertion follows from (17), since $Q/L \gg U/F_* \to \infty$. For the second assertion, let D consist of the captured denominators. By Theorem 3.4, $\operatorname{lcm}(D) < p$. Rational dissociation gives $2^{|D|}$ different rational values with reduced denominator at most p. Apply Theorem 3.1 with $K = 1$. □

This addresses the arithmetic exclusion step of a possible inverse Littlewood–Offord proof, rather than assuming a new random residue model. For example, a proof that a violation of the desired concentration bound forces a positive proportion of the true prime inverses into a GAP of rank $o(L/(U^2V))$ and formal volume at most $p\exp(-aL/(UV))$ would suffice. Standard polynomial-scale inverse Littlewood–Offord theorems do not give that conclusion at probability $p^{-1}\exp(aL/(UV))$; this is an exponential scale in $|B|$. The distinction is explicit in the classical inverse results [9, 10]. This manuscript does not extrapolate their constants to a growing exponent. Even a rank bound $O(L/U)$ and volume $p^{1-o(1)}$ is not sufficient

for (15). The obstruction does not establish that high concentration necessarily has such a representation.

3.5. **Uniform support expansion by distinct denominator copies.** The rank–height argument has an unconditional consequence for addition in $\mathbb{F}_p$. We use Green–Ruzsa's Freiman theorem [8, Theorem 1.1]: if a nonempty finite set S in an abelian group satisfies $|S+S| \le K_0|S|$, it is contained in a coset of a proper progression plus a subgroup, with rank $O(K_0^4 \log(K_0+2))$ and cardinality at most $|S| \exp(O(K_0^4 \log^2(K_0+2)))$. In the prime cyclic group the subgroup is trivial whenever the latter upper bound is less than p. For a proper progression its cardinality equals formal volume, so Theorem 3.4 applies.

Theorem 3.6 (Pointwise support expansion)**.** *Fix $0 < \gamma < 1/4$, $c_0 > 0$, and let $\tau = c_0 L/(UV)$. Let B consist of the primes in a fixed nonempty proportional subinterval of $[Q/2, Q]$, with Q in (3). For every set $S \subseteq \mathbb{F}_p$ containing all b^{-1}, $b \in B$, one has*

$$|S| \le p \exp(-\tau) \quad \Longrightarrow \quad |S+S| > L^{\gamma}|S|. \tag{18}$$

For $S = \Sigma_p(B)$ this also implies $|S+tS| > L^{\gamma/2}|S|$ for every nonzero $t \in \mathbb{F}_p$, when the displayed size condition holds.

Proof. Suppose $|S+S| \le L^{\gamma}|S|$ and $|S| \le pe^{-\tau}$. Green–Ruzsa yields a coset progression of cardinality at most

$$p \exp(-\tau + O(L^{4\gamma}U^2)) < p,$$

so the subgroup is trivial and the progression is proper. Its rank is $d = O(L^{4\gamma}U)$ and its formal volume $\mathcal{L}$ satisfies the same cardinality bound. Since $\gamma < 1/4$,

$$\log\left((d+2)!Q^{d+1}\mathcal{L}\right) \le L - \tau + O(L^{4\gamma}U^2) < L$$

for all sufficiently large p. Theorem 3.4 would imply $\operatorname{lcm}(B) < p$. This is impossible: the PNT gives $\log \operatorname{lcm}(B) \asymp Q$, and $Q/L \to \infty$. Thus (18) holds. For the second assertion, reciprocal subset sums satisfy $-S = S - \sum_{b\in B} b^{-1}$, so $|S+S| = |S-S|$. Ruzsa's triangle inequality gives $|S-S||S| \le |S+tS|^2$ because $t \ne 0$. This proves the claim. The triangle inequality follows directly by choosing one representation of each element of $S-S$ and injecting its product with the third set into the product of the two difference sets. □

For actual reciprocal subset sums one can implement the undilated addition in (18) without repeated denominators. For every $b \in B$ the identity

$$\frac{1}{b} = \frac{1}{2b} + \frac{1}{3b} + \frac{1}{6b} \tag{19}$$

gives a second copy of the binary choice, independent of the original choice of $1/b$: either take all three terms or none. Since b is a prime greater than 6, the sets $B, 2B, 3B, 6B$ are pairwise disjoint and have no internal coincidences. Thus

$$S+S \subseteq \Sigma_p(B \cup 2B \cup 3B \cup 6B), \qquad R(B \cup 2B \cup 3B \cup 6B) = 2R(B) = O(1/U).$$

In particular the growth bound is realized at denominator height $6Q$, with a controlled real cost, rather than by simply assuming repeated denominators are permitted.

Corollary 3.7 (A weaker support criterion with one extra copy)**.** *Use the larger-scale pointwise rectangular block from Section 2.1 and Proposition 2.3, writing $Q = q_p$ and $\tau = \tau_1 \asymp L/(UV)$ (changing $L = \log X$ to $\log p$ affects only constant bounds). Let B be its disjoint prime block in $[4Q_0/5, 9Q_0/10]$. For any fixed $0 < \gamma < 1/4$, the condition*

$$|\Sigma_p(B)| \ge p \exp(-4\tau + 2)L^{-\gamma} \tag{20}$$

already implies $g(p) \le 78Q$ for sufficiently large p. This relaxes the support threshold in Proposition 2.3 by a growing factor, although the logarithmic saving γU is $o(\tau)$.

Proof. If $|S| = |\Sigma_p(B)| \geq p\exp(-4\tau+2)$ use Proposition 2.3. Otherwise $|S| < pe^{-\tau}$, and Theorem 3.6 and (20) imply $|S+S| \geq pe^{-4\tau+2}$. Let ν be uniform on $S+S$ and let μ_A be the independent Bernoulli sum on the rectangular block A. As in Proposition 2.3, Parseval gives $\sum_h |\widehat{\nu}(h)|^2 = p/|S+S|$. For the nonzero frequencies with $E_A(h) \leq \tau$, use the established pointwise bound $\exp(O(U))$ for their number and the lower energy bound $\gg L/(U^2V)$; their total contribution to $\sum_h |\widehat{\mu}_A(h)\widehat{\nu}(h)|^2$ is $o(1)$. For the other frequencies the contribution is at most $e^{-4\tau}p/|S+S| \leq e^{-2}$. Thus the support of $\mu_A * \nu$ has size greater than $p/2$ by Cauchy–Schwarz. This argument uses uniformity on the set $S+S$, not a concentration assumption for the original Bernoulli weights.

Set $D = A \cup B \cup 2B \cup 3B \cup 6B$. The three new denominator blocks lie strictly above Q; all are disjoint, and A was disjoint from B. By (19), $\Sigma_p(D)$ contains the convolution support just considered. Since $D \subset [Q/2, 6Q]$, the sets D and $13D$ are disjoint. Their subset sum supports have cardinality greater than $p/2$, so $\Sigma_p(D) + 13^{-1}\Sigma_p(D) = \mathbb{F}_p$. The real cost is bounded by

$$R(D \cup 13D) = (1+1/13)(R(A)+2R(B)) \leq (14/13)(\log 2 + o(1)) < 1.$$

All denominators are at most $78Q < p$. This proves the claim with strict distinctness. □

The improvement (20) is an unconditional implication, not a proof that its hypothesis holds. In particular, even combining the sharp low-height entropy theorem with the expansion theorem gives only sublinear-in-L known lower bounds for the logarithm of the support. A bounded number of copies, needed to preserve an $O(LU)$ denominator bound, gains only $O(U)$ in the logarithm from the present quantitative doubling estimate. No near-uniform pointwise collision bound for the pure prime block has been proved here. A useful further question is whether the genuine reciprocal arithmetic forces substantially faster growth than (18), or whether an inverse concentration theorem can reach the rank–height budget of Corollary 3.5. Neither possibility is being used as a theorem.

4. Quasipolynomial covering obstruction and pointwise convolution flattening

The preceding rank–height theorem excludes a single sufficiently small progression containing all the reciprocal generators. We now exclude a union of many translates, even when it captures only a modest probability mass rather than the entire support. This yields an exponential-in-a-power-of-log p improvement to the doubling bound and, through Balog–Szemerédi–Gowers, a corresponding pointwise collision-probability *ratio* estimate for one additional copy. The result does not give the absolute concentration estimate at the original height.

Use the notation of Section 3: $L = \log p$, $U = \log L$, $V = \log U$, and Q in (3). Let B be all primes in a fixed nonempty proportional subinterval of $[Q/2, Q]$. Write $b_- = \min B$, $n = |B|$, $\mu = \mu_{p,B}$ for the law of its Bernoulli reciprocal subset sum, and $S_B = \Sigma_p(B)$. We use counting-measure norms $\|f\|_2^2 = \sum_{x\in\mathbb{F}_p} |f(x)|^2$, so $\rho_p(B) = \|\mu\|_2^2$. Uniformly in the specified range,

$$n \asymp Q/U = L^{1-o(1)}, \qquad n \gg L/\mathcal{C}_4(p), \qquad L/\log b_- = o(n), \qquad \log Q \sim U. \tag{21}$$

The last two assertions follow from $\mathcal{C}_4(p) = U^{o(1)}$ and the prime number theorem. The same arguments work for any subset of such primes satisfying these estimates. Constants in the following theorems may depend on the fixed proportional interval and the fixed constants in (3).

4.1. Lifting rational points and bounding the mass of a union of progressions.

Lemma 4.1 (Rational-height lifting)**.** *Let $\mathcal{P} \subseteq \mathbb{F}_p$ be an affine GAP of rank d and formal volume $\mathcal{L}$ (not necessarily proper). Let $\mathcal{R}$ be a finite nonempty set of rational numbers, each of reduced numerator and positive denominator of absolute value at most H, with H a positive integer. Assume all denominators are coprime to p, every member reduces into $\mathcal{P}$, and*

$$(d+2)!H^{d+2}\mathcal{L} < p. \tag{22}$$

Then there is a positive integer $T \le (d+1)!H^{d+1}\mathcal{L} < p$ divisible by all the reduced denominators in $\mathcal{R}$.

Proof. Write $v = a_v/b_v$ in reduced form, and choose one coordinate row for each v with leading coordinate one. The modular membership supplies the column-span hypothesis of Finite module 1. Now take $H_a = H_b = H$. Since $r \le d+1$, assumption (22) implies $(r+1)!H^{r+1}\mathcal{L} < p$. The module gives a positive common denominator $T_0 \le r!H^r\mathcal{L} \le (d+1)!H^{d+1}\mathcal{L}$. This proves the assertion for arbitrary rational inputs, including zero with reduced denominator one. □

For a GAP $\mathcal{P}$ of rank d and formal volume $\mathcal{L}$, the difference progression $\mathcal{P} - \mathcal{P}$ has a coordinate description with rank d and formal volume at most $2^d\mathcal{L}$, regardless of properness. Let $s \ge 1$ be an integer, and put

$$n_* = n - \lfloor L/\log b_- \rfloor, \qquad H = Q^{2s}.$$

By (21), $n_* \sim n$.

Theorem 4.2 (Progression-cover mass bound)**.** *Suppose $s \le n_*$ and*

$$(d+2)!\, Q^{2s(d+2)}\, 2^d\mathcal{L} < p. \tag{23}$$

For every set of J translates of $\mathcal{P}$, their union $\mathcal{U}$ satisfies

$$\mu(\mathcal{U}) \le \min\left(1, \frac{J}{\sum_{i=0}^{s}\binom{n_*}{i}}\right). \tag{24}$$

The same bound holds for the law of the independent three-valued differences with weights $1/2, 1/4, 1/4$.

Proof. Form the finite rational set consisting of 0 and all $v = \sum_{b\in D} e_b/b$ with $D \subseteq B$, $|D| \le 2s$, $e_b \in \{-2,-1,1,2\}$, and $v \pmod p \in \mathcal{P} - \mathcal{P}$. For nonempty D the reduced denominator is exactly $\prod_{b\in D} b$, since the numerator is nonzero modulo each $b \in D$. Also $|v| \le 2\sum_{b\in B} 1/b = O(1/U) < 1$, so its reduced numerator and denominator have height at most H. Apply Lemma 4.1 to this rational set and the difference progression. By (23), all its denominators divide an integer $T < p$. At most $\lfloor L/\log b_- \rfloor$ primes in B can divide T.

Fix all input coordinates for primes dividing T. On the remaining at least n_* coordinates, two distinct inputs whose modular sums lie in a particular translate of $\mathcal{P}$ cannot have Hamming distance at most $2s$. Otherwise their rational difference belongs to the preceding finite set, while its nonempty prime support is disjoint from the prime factors of T. This is impossible. Thus the preimage of the translate, on the unfixed coordinates, is a code of distance greater than $2s$. For binary inputs its radius-s Hamming balls are disjoint and each has probability mass at least $\sum_{i=0}^{s}\binom{n_*}{i}$ times the probability of its centre. For the three-valued law this follows as well: replacing any coordinate value c by all other values multiplies the combined coordinate probability by $(1-w(c))/w(c) \ge 1$. Summing over the disjoint balls yields the reciprocal ball-volume bound for the conditional probability. Average over the fixed coordinates, and sum over the J translates, to obtain (24). This proof does not replace a union of translates by a single progression of large rank. □

In particular, a covering of the entire subset-sum support by such translates must have $J \ge \sum_{i\le s}\binom{n_*}{i}$. More strongly, this is a quantitative arithmetic obstruction to concentrating a positive amount of the actual Bernoulli law inside a union of low-rank progressions of small volume. The sparse differences in the argument are used only to establish the common excluded prime support; the original subset sums need not be sparse.

4.2. **A quasipolynomial cover from small doubling.** We use two standard external additive-combinatorial theorems with their uniform quantitative statements. By Sanders's Bogolyubov–Ruzsa lemma [11, Theorem 1.1], if $D \subseteq \mathbb{F}_p$ is nonempty and $|D+D| \le K|D|$, then $2D-2D$ contains a proper symmetric coset progression P_0 of rank $O(\log^6(2K))$ and size at least

$$|D| \exp(-O(\log^6(2K) \log\log(4K))).$$

We have used the weaker size bound with $|D|$ rather than $|D+D|$. Plünnecke–Ruzsa gives $|jD-kD| \le K^{j+k}|D|$ for nonnegative integers j,k with $j+k>0$. If $K^4|D| < p$, the subgroup in the coset progression is trivial, so P_0 is an ordinary proper GAP of formal volume $|P_0| \le K^4|D|$. Choose a maximal collection of disjoint translates $x+P_0$ with $x \in D$. They lie in $D+P_0 \subseteq 3D-2D$, giving at most $K^5|D|/|P_0|$ such translates. Maximality covers D by the corresponding translates of P_0-P_0. Consequently there exists a GAP $\mathcal{P}$ and a covering of D by J translates of $\mathcal{P}$ with

$$\begin{aligned} d &\ll \log^6(2K), \\ \log J &\ll \log^6(2K)\log\log(4K) + \log(2K), \\ \mathcal{L} &\le 2^d K^4 |D|. \end{aligned} \tag{25}$$

Here $\mathcal{P} = P_0 - P_0$ may be nonproper, but the formal-volume bound holds in its coordinate description. This use of Sanders is different from applying a single-progression Freiman containment theorem, whose rank depends polynomially on K rather than on $\log K$. In the prime group the general abelian-group version of Sanders applies uniformly; a polynomial Freiman–Ruzsa theorem for bounded torsion with constants depending on that torsion would not give the needed uniform statement.

Fix $c_0 > 0$ and put

$$\tau = c_0 L/(UV), \qquad \lambda = c_1(\tau/U)^{1/12}, \tag{26}$$

where the fixed positive c_1 is sufficiently small. In particular $\lambda = L^{1/12-o(1)}$, $\lambda \gg U$, and $\lambda^{12}U = c_1^{12}\tau$. We record the parameter calculation used twice below. If a set is covered as in (25), with $\log(2K) \le C_2\lambda$ and $|D| \le p\exp(-\tau + O(\lambda))$, then, for p large,

$$d \le C_3\lambda^6, \quad \log J \le C_3\lambda^6 U, \quad \log \mathcal{L} \le L - \tau/2.$$

Take $s = \lceil D_0\lambda^6 \rceil$ with D_0 a sufficiently large fixed constant depending on C_3. Since $n_* = L^{1-o(1)}$ and $s = O(\sqrt{L})$, we have $\log(n_*/s) \ge U/3$ for sufficiently large p. Hence

$$\log \binom{n_*}{s} \ge s\log(n_*/s) \ge (D_0/3)\lambda^6 U.$$

On the other hand the logarithmic overhead in (23), besides $\log \mathcal{L}$, is

$$\log((d+2)!) + 2s(d+2)\log Q + d\log 2 = O_{C_3,D_0}(\lambda^{12}U).$$

First choose $D_0 > 3(C_3+1)$, then choose c_1 sufficiently small in (26) to make the last overhead at most $\tau/4$. Thus (23) holds, and the mass bound gives

$$\mu(D) \le \exp(-c_4\lambda^6 U) \tag{27}$$

with a fixed $c_4 > 0$. The constants are chosen before p tends to infinity, and are allowed to depend on the fixed C_2. Here and below $K^4|D| < p$ follows from the displayed hypotheses, so the use of an ordinary proper progression in (25) is justified.

Theorem 4.3 (Stronger uniform support expansion)**.** *With* (26) *and sufficiently small* c_1*, every set* $D \subseteq \mathbb{F}_p$ *containing* S_B *satisfies*

$$|D| \le pe^{-\tau} \quad \Longrightarrow \quad |D+D| > e^{\lambda}|D|. \tag{28}$$

For $D = S_B$*, also* $|S_B + tS_B| > e^{\lambda/2}|S_B|$ *for every nonzero* $t \in \mathbb{F}_p$ *under the size hypothesis.*

Proof. If $|D+D| \le e^{\lambda}|D|$, use (25) with $K = e^{\lambda}$ and the parameter calculation above. Formula (27) contradicts $\mu(D) = 1$. For S_B, reflection is a translate, so $|S_B - S_B| = |S_B + S_B|$. Ruzsa's triangle inequality gives $|S_B - S_B||S_B| \le |S_B + tS_B|^2$. This proves the second assertion. □

Unlike the earlier bound L^{γ} in Theorem 3.6, the gain in (28) exceeds every fixed power of L. The proof requires that D contain the full binary reciprocal subset-sum support, not just the singleton inverses. The gain for $t = 1/2$ is realized by the two disjoint denominator blocks $B, 2B$, at height $2Q$ with real cost $O(1/U)$.

4.3. **A pointwise collision-probability ratio theorem.** The mass result also applies to a moderately heavy subset of the law, not only its support. This enables a pointwise flattening estimate. We invoke the polynomial energy form of the Balog–Szemerédi–Gowers theorem (e.g. [13, Theorem 1.1 or 1.2] and Plünnecke–Ruzsa): a finite set D with additive energy $E^+(D) \ge |D|^3/K$ has a subset D' of size at least $cK^{-C}|D|$ and doubling $|D'+D'| \le CK^C|D'|$, with fixed positive constants. Only the polynomial form is needed.

Corollary 4.4 (One-copy pointwise flattening)**.** *Under* (21) *and* (26)*, for every sufficiently large prime p,*

$$p\rho_p(B) \ge e^{\tau} \quad \Longrightarrow \quad \|\mu * \mu\|_2^2 < e^{-\lambda}\rho_p(B). \tag{29}$$

In particular, with the same fixed constants,

$$\rho_p(B \cup 2B) \le \max\left(p^{-1}e^{\tau},\ e^{-\lambda}\rho_p(B)\right). \tag{30}$$

The second inequality concerns ordinary independent Bernoulli choices of mutually distinct denominators, not a model allowing repeated copies of $1/b$.

Proof. Theorem 5.3 below proves the same assertions, including the distinct-denominator realization, with the larger gain $\lambda_6 = L^{1/6-o(1)} > \lambda$ for all sufficiently large primes. The pure-prime mass estimate (27) also gives an independent route to this weaker rate via the dyadic-level/BSG argument used in that proof. We state this consequence to preserve the comparison with the sparse-difference lifting method, but do not duplicate the flattening proof. □

Corollary 4.5 (Two different prime blocks at a fixed height)**.** *Let B_1, B_2 be disjoint prime blocks to which Theorem 4.4 applies with common constants, for example two disjoint fixed proportional subintervals at the same scale Q. Write $\rho_i = \rho_p(B_i)$. If both $p\rho_i \ge e^{\tau}$, then*

$$\rho_p(B_1 \cup B_2) \le e^{-\lambda}\sqrt{\rho_1\rho_2}.$$

If either $p\rho_i < e^{\tau}$, then $p\rho_p(B_1 \cup B_2) < e^{\tau}$. Thus the ratio estimate can also be used for two independently chosen prime blocks without increasing the denominator height.

Proof. This is the mixed-convolution conclusion of Theorem 5.3, with the weaker rate λ. Its proof is included there. □

4.4. **A quantitative relaxation of the covering criterion.** At the larger scale $Q = q_p = \lfloor C \log p \log_2 p \rfloor$ use the prime block B in $[4Q_0/5, 9Q_0/10]$ and the disjoint rectangular block A in Section 2.1. Use $\tau = \tau_1$ of the pointwise spectral input in Section 2.1, and λ of (26) (the dyadic and prime logarithms are comparable and the proof applies with τ_1 itself).

Corollary 4.6 (Coverage after one distinct dilated copy)**.** *Each of the following two conditions at a given sufficiently large prime implies $g(p) \le 10Q$:*

$$|\Sigma_p(B)| \ge p\exp(-4\tau + 2 - \lambda/2),$$
$$\log(p\rho_p(B)) \le 4\tau - 2 + \lambda.$$

Proof. Under the first condition, if $|S_B| \ge pe^{-4\tau+2}$ the enlarged support already meets the threshold. Otherwise $|S_B| < pe^{-\tau}$ and Theorem 4.3 with $t = 2^{-1}$ shows that $|\Sigma_p(B \cup 2B)| \ge pe^{-4\tau+2}$. Under the second condition, Theorem 4.4 gives $\rho_p(B \cup 2B) \le p^{-1}e^{4\tau-2}$ (since $e^{\tau} \le e^{4\tau-2}$ for large p). Cauchy–Schwarz gives the same support threshold. In both cases let

ν be uniform on that enlarged support. Apply the Fourier proof of Proposition 2.3 unchanged to the convolution of ν and the Bernoulli sum on A: the low-energy squared Fourier mass is $o(1)$ and the high-energy squared mass is at most $e^{-4\tau}p/|\Sigma_p(B \cup 2B)| \le e^{-2}$. Thus the support of $D = A \cup B \cup 2B$ has size greater than $p/2$. All the denominator sets are disjoint, because A consists of composites at most Q, whereas B consists of primes and $2B$ lies above Q. Also $D \subseteq [Q/2, 2Q]$ and $R(D) = R(A) + (3/2)R(B) = o(1)$, by the sizes of the blocks and their minimum denominators. The set $5D$ lies above $2Q$ and is disjoint from D. Since two subsets of $\mathbb{F}_p$ of size greater than $p/2$ have full sumset, $\Sigma_p(D) + 5^{-1}\Sigma_p(D) = \mathbb{F}_p$. The total real mass $(6/5)R(D) = o(1) < 1$ and all denominators are at most $10Q < p$. This proves the corollary. □

Remark 4.7 (What the new pointwise result does and does not establish). The exponential improvement $\lambda = c_1(\tau/U)^{1/12} = L^{1/12-o(1)}$ is a genuine gain over the previous $O(U)$ logarithmic support increment. The flattening theorem is a ratio bound for the collision probability after adding a disjoint scaled block, not an absolute $p^{-1}\exp(o(\tau))$ bound for the original prime block. Both conditions of Corollary 4.6 remain unproved uniformly; in particular $\lambda = o(\tau)$, and the available absolute collision-entropy lower bounds are still $o(L)$. Repeated doubling does not solve this at fixed height: to force an $O(L)$ entropy gain from these bounds could require order L/λ doublings, and the needed independent copies cannot simply be placed within a constant-factor denominator budget. The paper does not claim a new all-primes height upper bound from this section. The argument suggests a specific bottleneck for a structural method: the logarithmic covering number and the rank of the progression both enter the arithmetic height budget. Stronger information about the actual reciprocal distribution, rather than only general small-doubling structure, is still needed to attain the target exponent. The proofs in this section have been checked internally and depend on the stated external BSG, Sanders and Plünnecke–Ruzsa theorems; no claim of priority or third-party review is made.

5. Additive complexity budget from Boolean-cube sumsets

The multiplicative overhead in (23) can be avoided. We use the Boolean-cube sumset inequality of Hajela–Seymour and Landau–Logan–Shepp [14, 15]; see also [16, Corollary 2 and historical discussion] for the more general discrete-cube inequality. For a nonempty $E \subseteq \{0,1\}^t$ and integer $k \ge 1$ it gives, in $\mathbb{Z}^t$,

$$|kE| \ge |E|^{\log(k+1)/\log 2}. \tag{31}$$

The exponent is uniform in both t and k; the notation kE denotes a k-fold sumset. Only this classical Boolean case is used. The new argument conditions on a fixed set of prime coordinates and allows controlled modular wraparound of its expanded digit box. In particular, there is no need to lift the generators of a GAP to rational numbers or to clear distinct denominators for each row of an augmented matrix.

We retain the prime block B, its law $\mu = \mu_{p,B}$, and the notation of (21). In particular $L = \log p$, $U = \log L$, $V = \log U$, $\log Q \sim U$, and $n = |B| \gg L/\mathcal{C}_4(p)$ with $\mathcal{C}_4(p) = U^{o(1)}$.

Finite module 2 (Integer digit injection and modular folding)**.** *Let R be a nonempty finite set of distinct primes, $P_R = \prod_{b\in R} b$, $t = |R|$, and $0 \le k < \min R$ an integer. For $z \in \{0,\dots,k\}^R$ put*

$$N(z) = \sum_{b\in R} (P_R/b) z_b \in \mathbb{Z}.$$

Then N is injective and $0 \le N(z) < tP_R$. For any prime $p \nmid P_R$, each residue of $N \bmod p$ has at most $1 + \lfloor tP_R/p \rfloor$ digit preimages. In particular, if $P_R < p$ the bound $t+1$ suffices. The reciprocal image of a digit set Z modulo p consequently has cardinality at least $|Z|/(t+1)$ when $P_R < p$.

Proof. Modulo b, all terms in $N(z)$ except the b-term vanish, and P_R/b is a unit. Therefore equality of integer numerators implies $z_b = z'_b \pmod b$; the digit bounds force equality. The

range follows from $z_b/b < 1$ for each b. Any two distinct integers with the same residue modulo p differ by at least p, so the stated range yields the fibre bound. Finally P_R is a unit modulo p, and multiplying the reciprocal image by it gives $N \bmod p$. These arguments do not require the modular digit map to be injective. □

Theorem 5.1 (Additive progression-cover budget)**.** *There exist a set $R \subset B$, $t = |R| \asymp L/U$, and an integer $k \asymp Q$ with $1 \le k < \min R$ such that, with*

$$\beta = \frac{\log(k+1)}{\log 2}, \qquad E_0 = L - t\log(k+1) + \log(t+1),$$

one has $0 \le E_0 = O(L/U^{3/2} + U)$. For every union $\mathcal{U}$ of $J \ge 1$ translates of a GAP of rank d and formal volume $\mathcal{L}$, and every probability measure ν on $\mathbb{F}_p$,

$$(\mu * \nu)(\mathcal{U}) \le \min\left\{1, J\exp\left(\frac{\log\mathcal{L} - L + d\log k + E_0}{\beta}\right)\right\}. \tag{32}$$

In particular, if $\delta = (\mu*\nu)(\mathcal{U}) > 0$, then

$$L \le \log\mathcal{L} + d\log k + \beta\log J + \beta\log(1/\delta) + E_0. \tag{33}$$

The constants are uniform for the fixed proportional prime blocks in (21). Properness of the progression is not required.

Proof. Partition the fixed proportional prime interval into $O(\sqrt{U})$ intervals whose endpoint ratios are at most $1 + U^{-1/2}$. One contains $\gg n/\sqrt{U} \gg L/U$ primes, since $\mathcal{C}_4(p) = U^{o(1)}$. Its prime product exceeds p. Choose primes successively from that interval until multiplying by the next would exceed p, and let R consist of the selected primes. Its product P_R satisfies $p/Q < P_R < p$ (equality with p is impossible). Thus $t \asymp L/U$. Put $k = \min R - 1$. All $b \in R$ satisfy $b/(k+1) \le 1 + U^{-1/2}$, so

$$0 \le L - t\log(k+1) = \log(p/P_R) + \sum_{b\in R}\log(b/(k+1)) \ll U + L/U^{3/2}.$$

Also $\log(t+1) = O(U)$.

Consider a single translate $\mathcal{P}$ of the given GAP with $\mu(\mathcal{P}) = \delta_0 > 0$. Conditional on the Bernoulli coordinates in $B \setminus R$, there is a choice of those coordinates for which at least $\delta_0 2^t$ vectors in $\{0,1\}^t$ have their reciprocal image in a translate of $\mathcal{P}$. Denote this set of vectors by E. Formula (31) yields

$$|kE| \ge (k+1)^t\delta_0^\beta.$$

For each $z \in kE$, the coordinates satisfy $0 \le z_b \le k < b$. Finite module 2 gives at least $|kE|/(t+1)$ distinct reciprocal residues, since $P_R < p$. On the other hand it lies in a translate of $k\mathcal{P}$, which has at most $k^d\mathcal{L}$ elements by its coordinate description. Consequently

$$(k+1)^t\delta_0^\beta/(t+1) \le k^d\mathcal{L}.$$

This is exactly (32) for $J = 1$ and ν a point mass, with the same constants for every translation. The union bound gives it for J translates. Finally $(\mu*\nu)(\mathcal{U}) = \sum_y \nu(y)\mu(\mathcal{U} - y)$, and the preceding uniform bound applies termwise. This proves the assertion, including (33). In particular the three-valued difference law of two Bernoulli sums is covered by taking ν to be the reflected Bernoulli law. □

The term for the rank in (33) is $d\log k$, the term for the logarithmic covering number is $\beta\log J$, and the term for the mass deficit is $\beta\log(1/\delta)$. They are additive. The previous product $2s(d+2)\log Q$ in (23) is not needed. The error E_0 is $o(\tau)$ for $\tau = c_0L/(UV)$, because $V/\sqrt{U} \to 0$. The large k used in this proof is an auxiliary integer sumset order, not a permitted multiplicity of the final unit-fraction expansion.

5.1. **Improved pointwise expansion and robust convolution flattening.** Fix $c_0 > 0$ and put

$$\tau = c_0 L/(UV), \qquad \lambda_6 = c_6(\tau/U^2)^{1/6}, \tag{34}$$

with a sufficiently small fixed $c_6 > 0$. Then $\lambda_6 = L^{1/6-o(1)}$, $\lambda_6 \gg U$, and $\lambda_6 = o(\tau/U)$. Use the quasipolynomial cover (25), which follows from the general-abelian-group Bogolyubov–Ruzsa theorem of Sanders [11] and Plünnecke–Ruzsa. If a nonempty set D satisfies

$$|D| \le p\exp(-\tau + C_7\lambda_6), \qquad |D+D| \le K|D|, \qquad \log(2K) \le C_7\lambda_6,$$

where C_7 is any fixed constant, then $K^4|D| < p$. Its covering parameters in (25) satisfy

$$d = O(\lambda_6^6), \qquad \log J = O(\lambda_6^6 U), \qquad \log \mathcal{L} \le L - \tau + O(\lambda_6 + \lambda_6^6).$$

All implied constants may depend on C_7. Inserting these inequalities into (33), and using $\beta \asymp U$, gives, for any probability ν,

$$(\mu * \nu)(D) \le \exp(-c\tau/U). \tag{35}$$

Indeed the rank and logarithmic covering costs together are $O(\lambda_6^6 U^2) = O(c_6^6\tau)$. Choose c_6 small enough, depending on C_7, to make them at most $\tau/4$, and absorb $E_0 = o(\tau)$ and $O(\lambda_6 + \lambda_6^6)$ in the remaining margin. This proves (35) with a fixed $c > 0$. In particular the gain in the exponent has changed from $1/12$ to $1/6$ without a change to the external Sanders theorem.

Theorem 5.2 (Pointwise support expansion at the improved rate)**.** *For every $D \subseteq \mathbb{F}_p$ containing a translate of $S_B = \Sigma_p(B)$,*

$$|D| \le pe^{-\tau} \quad \Longrightarrow \quad |D+D| > e^{\lambda_6}|D|.$$

If $|S_B| \le pe^{-\tau}$, then for every $a \ne 0$ in $\mathbb{F}_p$,

$$|S_B + aS_B| > e^{\lambda_6/2}|S_B|.$$

In particular the latter estimate is achieved using the distinct denominators in $B \cup 2B$ when $a = 2^{-1}$.

Proof. A translate of μ has mass one on D. If the doubling were at most e^{λ_6}, (35) would give a contradiction. Reflection of the binary subset-sum support S_B is a translate of S_B (complement every subset), so $|S_B - S_B| = |S_B + S_B|$. Ruzsa's triangle inequality gives $|S_B - S_B||S_B| \le |S_B + aS_B|^2$. This proves the second estimate. The prime and doubled-prime blocks are disjoint for large p. □

Finite norm interface for the flattening theorems. The energy extraction and the structural theorem are logically separate. In the next module all norms on the finite group are counting-measure norms and the Fourier L^4 norm is taken with normalized counting measure on its dual. No asymptotic assumption on p is made.

Finite module 3 (Dyadic energy extraction with explicit losses)**.** *Let G be a finite abelian group of size $P \ge 2$ and η a probability law, with $\rho = \|\eta\|_2^2$. Let $0 < \xi \le 1$ satisfy $2P^{-2} \le \sqrt{\xi}/2$ and suppose $\|\eta * \eta\|_2^2 \ge \xi\rho$. Write $J_* = 1 + \lceil 3\log P/\log 2\rceil$ and $K_1 = 64J_*^4/\xi$. Then there is a nonempty set D on which $w \le \eta(x) < 2w$ for some $w > 0$, with*

$$w^4E^+(D) \ge \rho/K_1, \quad E^+(D) \ge |D|^3/K_1, \quad \eta(D) \ge K_1^{-1/2}, \quad |D| \le K_1/\rho. \tag{36}$$

Proof. Put $f = \eta\mathbf{1}_{\eta \ge P^{-3}}$ and $g = \eta - f$. The discarded mass is at most P^{-2}, so Young's inequality gives $\|\eta*\eta - f*f\|_2 \le 2P^{-2}\sqrt{\rho}$. Hence $\|f*f\|_2 \ge \sqrt{\xi\rho}/2$. Divide f into the nonempty dyadic layers f_j defined by $2^{-j} \le \eta < 2^{1-j}$ for integers $j \ge 0$, including the layer $j = 0$ for the atom 1. There are at most J_* layers. Parseval, the convolution identity and Minkowski give

$$\|f*f\|_2^{1/2} = \|\widehat{f}\|_4 \le \sum_j \|\widehat{f_j}\|_4 \le J_* \max_j \|f_j * f_j\|_2^{1/2}.$$

Some layer satisfies $\|f_j * f_j\|_2^2 \ge \xi\rho/(4J_*^4)$. Its expansion as a weighted additive-energy sum is at most $16w^4E^+(D)$, giving the first inequality. With $A = w|D|$ and $B = w^2|D|$, we have $A \le 1$, $B \le \rho$, and $\rho/K_1 \le w^4|D|^3 = A^2B \le \rho$ because $E^+(D) \le |D|^3$. These give the energy bound, $A \ge K_1^{-1/2}$, and $B \ge \rho/K_1$. Since $|D| = A^2/B$, the size bound follows. □

Finite module 4 (Power interpolation for finite convolution)**.** *For probability laws μ, ζ on a finite abelian group of size P and any integer $j \ge 1$, with $\rho = \|\zeta\|_2^2$, one has*

$$\left(\frac{\|\zeta * \mu\|_2^2}{\rho}\right)^j \le \frac{\|\zeta * \mu^{*j}\|_2^2}{\rho} \le \frac{\|\mu^{*j}\|_\infty}{\rho}. \tag{37}$$

If $\|\mu^{*j}\|_\infty \le P^{-1}e^E$, then $d(\zeta * \mu) \le \min(d(\zeta), (1-1/j)d(\zeta) + E/j)$, where $d(\zeta) = \log(P\rho)$. In addition

$$\|f * g\|_2^2 \le \sqrt{\|f * f\|_2^2 \|g * g\|_2^2} \tag{38}$$

for any two real or complex functions on the group.

Proof. Parseval makes $|\widehat{\zeta}(h)|^2/(P\rho)$ a probability law on the characters. Apply Jensen's inequality for the integer power j to $|\widehat{\mu}(h)|^2$. Parseval and the convolution identity give the first inequality. The convolved law has mass one and supremum at most $\|\mu^{*j}\|_\infty$, giving the second. Young's inequality gives the monotonicity term; take logarithms of the first bound to obtain the other term. Finally Fourier Cauchy–Schwarz gives (38) with exactly the stated normalization. □

Theorem 5.3 (Robust pointwise flattening)**.** *Let ν be any probability measure on $\mathbb{F}_p$, and put $\eta = \mu * \nu$ and $\rho_\eta = \|\eta\|_2^2$. If $p\rho_\eta \ge e^\tau$, then*

$$\|\eta * \eta\|_2^2 < e^{-\lambda_6}\rho_\eta. \tag{39}$$

Consequently

$$\rho_p(B \cup 2B) \le \max(p^{-1}e^\tau, e^{-\lambda_6}\rho_p(B)). \tag{40}$$

For the Bernoulli distribution on any two disjoint prime blocks B_1, B_2 satisfying (21) with common constants, if $p\rho_p(B_i) \ge e^\tau$ for both i, then

$$\rho_p(B_1 \cup B_2) \le e^{-\lambda_6}\sqrt{\rho_p(B_1)\rho_p(B_2)}.$$

If either input is below the threshold, the union is also below the threshold. The stronger geometric-mean bound is not asserted in this latter case.

Proof. If (39) fails, apply Finite module 3 with $P = p$ and $\xi = e^{-\lambda_6}$. Here $\lambda_6 = o(L)$, so the truncation hypothesis holds; $\lambda_6 \gg U$ gives $\log K_1 = \log 64 + 4\log J_* + \lambda_6 = O(\lambda_6)$ with an absolute implied constant. The module supplies the set D, the energy bound, $\eta(D) \ge w|D| \ge K_1^{-1/2}$ and $|D| \le K_1/\rho_\eta$. The polynomial Balog–Szemerédi–Gowers theorem (as used in Theorem 4.4, see [13]) yields a subset $D' \subseteq D$ of mass at least $e^{-O(\lambda_6)}$ for η and doubling at most $e^{O(\lambda_6)}$. Its size is at most $K_1/\rho_\eta \le p\exp(-\tau + O(\lambda_6))$. Fix C_7 in (35) to accommodate the BSG losses, and choose c_6 accordingly. Formula (35) contradicts this lower bound for $\eta(D')$, because $\tau/U \gg \lambda_6$. This proves (39).

Take ν to be a point mass at zero. If $p\rho_p(B) \ge e^\tau$, Fourier Cauchy–Schwarz gives $\|\mu * \mu_{2^{-1}}\|_2^2 \le \|\mu * \mu\|_2^2$, and the left side is $\rho_p(B \cup 2B)$. If $p\rho_p(B) < e^\tau$, use Young's inequality instead. This gives (40). For two different blocks apply Fourier Cauchy–Schwarz to their individual self-convolution estimates if both are above the threshold, and Young's inequality otherwise. □

Corollary 5.4 (Relaxed sufficient condition at the larger height)**.** *At $Q \asymp LU$ use the disjoint rectangular block and the fixed prime block of Corollary 4.6, and take $\tau = \tau_1$ there. Each of the conditions*

$$|S_B| \ge pe^{-4\tau+2-\lambda_6/2}, \qquad \log(p\rho_p(B)) \le 4\tau - 2 + \lambda_6$$

separately suffices for $g(p) \le 10Q$.

Proof. Theorem 5.2 or Theorem 5.3, with the below-threshold alternative if necessary, gives $|\Sigma_p(B \cup 2B)| \geq pe^{-4\tau+2}$. The smoothing and disjoint dilation argument of Corollary 4.6 applies without change. In particular the sumset after adding the rectangular block has more than $p/2$ elements, and a disjoint fivefold dilation completes the coverage. The total real cost is $o(1)$ and all denominators are distinct. □

Remark 5.5 (Scope of the improvement). The additive budget is a direct obstruction for the actual reciprocal law, rather than a new general inverse theorem for small-doubling sets. It uses an auxiliary k-fold sumset with $k \asymp Q$ and hence must not be read as a coverage construction with that many repeated denominators. The passage from a self-convolution to $B \cup 2B$ in (40) uses independent, genuinely distinct denominators. For the more general law $\eta = \mu * \nu$, (39) does not assert that adding just one more copy of B necessarily flattens η by the same factor. The absolute pointwise collision-entropy gap remains: $\lambda_6 = o(L)$, and no bound in this section implies (65). At the sharp counting scale a block of primes alone has $n = o(L)$ and cannot cover $\mathbb{F}_p$; the improved relative expansion there is not an absolute coverage result either. If a future uniform structural cover has parameters $d(K)$ and $J(K)$, the new arithmetic obstruction charges $O(U(d(K) + \log J(K)))$, together with its formal-volume loss. This is the precise additive complexity budget that replaces the product in the previous section. The new arguments rely on the cited classical Boolean-cube inequality and on the stated external Sanders, BSG and Plünnecke–Ruzsa theorems. No new uniform bound for $f(p)$, $g(p)$, or $D(b)$ is claimed here.

6. Hierarchical digit boxes and heterogeneous convolution

We now use a composite-denominator block to replace the large auxiliary sumset order in Section 5 by a polylogarithmic order. The resulting mass inequality is valid for arbitrary sets rather than only progression covers. Plünnecke's inequality then suffices for support growth and, together with polynomial Balog–Szemerédi–Gowers, for symmetric L^2 flattening. A second, binomial-digit construction gives a genuinely asymmetric convolution estimate with arbitrary input noise. All these estimates are pointwise. The near-full digit capacity in this section is at the larger height $O(LU)$, not at the counting-optimal height $O(LU/\mathcal{C}_4(p))$; the final subsection proves an obstruction to transferring the argument at bounded common-denominator height.

6.1. **A rectangular block with its outer primes.** Fix $\epsilon > 0$, put $L = \log p$, $U = \log L$, $V = \log U$, and $\tau = c_0 L/(UV)$ for a fixed $c_0 > 0$. Constants may depend on ϵ and c_0. Choose a fixed $a > 4/\min(\epsilon, 1)$, let S be the primes in $(aU, 2aU]$, and select a set R of primes in $(2L, 4L]$ with

$$p/(4L) < P := \prod_{r \in R} r < p. \tag{41}$$

This is possible by the prime number theorem and a greedy choice. Define

$$A = R \cup \{rs : r \in R,\ s \in S\}, \quad T = \prod_{s \in S} s, \quad D = PT.$$

The denominators are distinct, at most $Q = \lceil CLU \rceil$ for $C > 8a$, and

$$|R| \asymp L/U, \quad |S| \sim aU/V, \quad \log T \sim aU, \quad |A| \asymp_\epsilon L/V, \quad \sum_{d \in A} d^{-1} = O_\epsilon(U^{-1}). \tag{42}$$

Write $\mu = \mu_{p,A}$ for its Bernoulli law, and $S_A = \operatorname{supp} \mu = \Sigma_p(A)$. One may also take R in another fixed interval $(cL, (c+2)L]$ and choose its product as in (41), with $4L$ replaced by $(c+2)L$. This provides any fixed number of disjoint copies at a constant-factor larger height, with the same S if desired.

Lemma 6.1 (A digit-box fiber bound). *Put $k = \lceil U^{1+\epsilon} \rceil$ and let the ξ_d, $d \in A$, be independent uniform variables on $\{0, \dots, k\}$. For some $E_0 \geq 0$ with*

$$E_0 = O_\epsilon\left(U + \frac{L}{U^{1+\epsilon}V}\right) = o(\tau), \tag{43}$$

one has

$$\max_{x\in\mathbb{F}_p}\mathbb{P}\left(\sum_{d\in A}\xi_d d^{-1}=x \pmod p\right)\le p^{-1}e^{E_0}. \tag{44}$$

Proof. Fix $r\in R$, and write $W_r=\xi_r+\sum_{s\in S}\xi_{rs}/s$. Partition each inner digit interval into complete blocks of length s and a possible final incomplete block. The probability θ_s of the incomplete block is at most $s/(k+1)$. Condition on the selected blocks and on the values of the incomplete coordinates. For the remaining complete coordinates S', the remainders u_s are independent and uniform modulo s. Hence, by CRT, their reciprocal sum is $v/T'+j(v)$, where $T'=\prod_{s\in S'}s$, v is uniform on $[0,T'-1]\cap\mathbb{Z}$, and $0\le j(v)\le|S|$. The case $T'=1$ uses $v=j(v)=0$. The remaining parts of the original inner digits contribute a fixed rational translate with denominator coprime to r.

After removing the translate, a local equality modulo r takes the form

$$v+T'z\equiv T'x \pmod r,\qquad z=\xi_r+j(v)\in[0,k+|S|].$$

For each integer z, at most $T'/r+1$ values of v satisfy the congruence, and (v,z) determines ξ_r uniquely. Thus the conditional local point probability is at most

$$\left(1+\frac{|S|}{k+1}\right)(r^{-1}+(T')^{-1}).$$

The indicators of incompleteness are independent. Therefore

$$\mathbb{E}(T')^{-1}=\prod_{s\in S}\left(\theta_s+\frac{1-\theta_s}{s}\right)\le\prod_{s\in S}\left(\frac{s}{k+1}+\frac{1}{s}\right)\le\exp\big(-[a\min(\epsilon,1)+o(1)]U\big)\ll L^{-3}.$$

It follows that the maximal local point probability of W_r modulo r is

$$r^{-1}\left(1+O_\epsilon(U^{-\epsilon}/V)\right). \tag{45}$$

The integer numerator $N=D\sum_{d\in A}\xi_d/d$ satisfies $N=\sum_{r\in R}(P/r)TW_r$. Its value fixes $W_r \pmod r$ for every r, since $(P/r)T$ is invertible modulo r. Independence of the rows yields the integer point bound

$$P^{-1}\exp\Big(O_\epsilon(L/(U^{1+\epsilon}V))\Big).$$

The full integer range is contained in $[0,O_\epsilon(Dk/U)]$ by (42). A residue modulo p therefore has at most $1+O_\epsilon(Dk/(pU))$ possible integer representatives in that range. Its logarithm is $O_\epsilon(U)$, since $D/p<T$ and $\log T=O_\epsilon(U)$. Together with (41) this proves (44) and (43). Reduction of the rational sum is legitimate because $(D,p)=1$. □

6.2. **An arbitrary-set mass inequality and asymmetric support growth.**

Theorem 6.2 (Hierarchical cube mass inequality)**.** *Set $\beta=\log(k+1)/\log 2\asymp_\epsilon V$, and let $\eta=\mu*\nu$ for any probability measure ν on $\mathbb{F}_p$. For every nonempty $Z\subseteq\mathbb{F}_p$,*

$$|kZ|\ge pe^{-E_0}\eta(Z)^\beta. \tag{46}$$

If $\eta(Z)>0$ and $|Z+Z|\le K|Z|$, then

$$\log(p/|Z|)\le E_0+k\log K+\beta\log(1/\eta(Z)). \tag{47}$$

Proof. First suppose ν is a point mass at zero, and let E be the preimage of Z under the reciprocal map of the Boolean cube $\{0,1\}^{|A|}$. Then $|E|=2^{|A|}\mu(Z)$. The classical inequality (31) gives $|kE|\ge(k+1)^{|A|}\mu(Z)^\beta$. Every fiber of the reciprocal map on the expanded digit box has size at most $(k+1)^{|A|}p^{-1}e^{E_0}$ by Lemma 6.1. Since the image of kE lies in kZ, this proves (46). The assertion is trivial when E is empty. In general some translate $Z-y$ has $\mu(Z-y)\ge\eta(Z)$, and $|k(Z-y)|=|kZ|$. For (47) apply the standard Plünnecke bound $|kZ|\le K^k|Z|$ (see, e.g., [12]). No progression containment is needed. □

When only support matters, a smaller sumset order suffices.

Proposition 6.3 (Asymmetric pointwise support growth)**.** *Set $k_0 = \lceil 2aU \rceil$. For every nonempty $Z \subseteq \mathbb{F}_p$ one has*

$$\log \frac{|Z + S_A|}{|Z|} \geq \frac{\log(p/|Z|) - O_\epsilon(U)}{k_0}. \tag{48}$$

In particular, when $|Z| \leq pe^{-\tau}$, adding a single distinct block A to a disjoint set of denominators with reciprocal support Z increases that support by at least $\exp(c_\epsilon L/(U^2V))$.

Proof. In each row choose inner digits $c_s \in [0, s-1]$. Their sum is $v/T + j(v)$, $0 \leq j(v) \leq |S|$, with v ranging over all $[0, T-1] \cap \mathbb{Z}$. The outer digit $c_r = |S| - j(v) \in [0, k_0]$ gives the value $|S| + v/T$. Since $T > r$, this covers every local residue modulo r. Choosing independently in every row yields at least P distinct integer numerators modulo P, after multiplication by D. The possible integer numerators lie in $[0, O_\epsilon(Dk_0/U)]$. Thus their distinct values modulo p number at least $P/(1 + O_\epsilon(Dk_0/(pU))) \geq pe^{-O_\epsilon(U)}$. All digits lie in $[0, k_0]$, so these images belong to $k_0 S_A$. Now if $K = |Z + S_A|/|Z|$, the Plünnecke bound $|k_0 S_A| \leq K^{k_0}|Z|$ proves (48). The last assertion follows since $U = o(\tau)$ and $k_0 \asymp_\epsilon U$. The high-order auxiliary sumset is not a repeated-denominator representation in the conclusion; $Z + S_A$ uses the binary block A just once. □

6.3. **Symmetric flattening without a progression-cover theorem.**

Theorem 6.4 (Pointwise symmetric flattening from the cube budget)**.** *Let $\eta = \mu * \nu$, with ν any independent probability law, and write $\rho_\eta = \|\eta\|_2^2$. There is a sufficiently small constant $c_1 > 0$ such that, with*

$$\lambda_7 = c_1\tau/k \asymp_\epsilon L/(U^{2+\epsilon}V), \tag{49}$$

one has

$$p\rho_\eta \geq e^\tau \quad \Longrightarrow \quad \|\eta * \eta\|_2^2 < e^{-\lambda_7}\rho_\eta.$$

For two independent laws η_1, η_2 each containing such an anchor block (possibly different) and with both $p\rho_{\eta_i} \geq e^\tau$, their mixed convolution has squared L^2 norm at most $e^{-\lambda_7}\sqrt{\rho_{\eta_1}\rho_{\eta_2}}$, using common constants. If either input is below the threshold, the convolution is also below that threshold.

Proof. Suppose the assertion fails. Apply Finite module 3 with $P = p$ and $\xi = e^{-\lambda_7}$. The truncation hypothesis holds since $\lambda_7 = o(L)$; moreover $\lambda_7 \gg U$, so $\log K_1 = O(\lambda_7)$. With $Z = D$ from the module, the conclusions of (36) give

$$E^+(Z) \geq |Z|^3/K_1, \qquad w|Z| \geq K_1^{-1/2}, \qquad |Z| \leq K_1/\rho_\eta.$$

The polynomial BSG theorem gives $Z' \subseteq Z$ with $|Z'| \geq K_1^{-O(1)}|Z|$ and $|Z' + Z'| \leq K_1^{O(1)}|Z'|$. For example the difference-set form in [13], followed by Plünnecke–Ruzsa, suffices. Hence $\log(1/\eta(Z')) = O(\lambda_7)$ and $\log(p/|Z'|) \geq \tau - O(\lambda_7)$. Inequality (47) now implies

$$\tau - O(\lambda_7) \leq E_0 + O(k\lambda_7 + \beta\lambda_7).$$

Since $E_0 = o(\tau)$ and $\beta \ll k$, choosing c_1 sufficiently small gives a contradiction. Finally Fourier Cauchy–Schwarz bounds the squared L^2 norm of a mixed convolution by the geometric mean of the two squared self-convolution norms. If either input has small squared L^2 norm use Young's inequality. All statements are uniform for fixed anchor parameters and independent of ν. □

The rate (49) is $L^{1-o(1)}$ and strictly improves the earlier $L^{1/6-o(1)}$ one-copy logarithmic gain, for these composite-anchor laws at the larger scale LU. To turn self-convolution into distinct denominators, Fourier Cauchy–Schwarz also bounds $\rho(\eta * (2^{-1}\eta))$ by $\|\eta * \eta\|_2^2$, where $2^{-1}\eta$ denotes the image law under multiplication by 2^{-1}. If η is the Bernoulli law of a set D_1 with $D_1 \cap 2D_1 = \varnothing$, this mixed convolution uses precisely $D_1 \cup 2D_1$. In particular A is odd and has this property. For arbitrary additional denominators one must check disjointness; no automatic disjointness of D_1 and its dilate is assumed.

6.4. **A binomial-digit method for genuine asymmetric anticoncentration.** The preceding symmetric theorem does not assert the same ratio bound for adding A once to a law of arbitrary, different entropy. A direct Fourier power argument gives a separate asymmetric result. Put $\ell = \lceil U^{2+2\epsilon} \rceil$.

Lemma 6.5 (Binomial high power of the anchor law)**.** *One has*

$$\|\mu^{*\ell}\|_\infty \le p^{-1} e^{E_1}, \qquad E_1 = O_\epsilon\left(U + \frac{L}{U^{1+\epsilon}V}\right) = o(\tau), \tag{50}$$

where $E_1 \ge 0$. Here the repeated sums are auxiliary random variables.

Proof. The law $\mu^{*\ell}$ uses independent binomial digits, each of law $w(j) = 2^{-\ell}\binom{\ell}{j}$ on $0 \le j \le \ell$ and zero elsewhere. This probability sequence is unimodal and $\|w\|_\infty = O(\ell^{-1/2})$. Partition the integers into consecutive blocks I of length s. On each block subtract the uniform submeasure of point weight $\min_I w$. These submeasures, together with a nonnegative residual measure, form a mixture decomposition of w. The residual mass is bounded by

$$\theta_s \le s \sum_I (\max_I w - \min_I w) \le s \sum_{j \in \mathbb{Z}} |w(j+1) - w(j)| = 2s\|w\|_\infty.$$

Choose independently whether each inner digit comes from a uniform block or the residual part, and condition on the chosen uniform blocks and the actual residual digits. Exactly as in Lemma 6.1, the good-coordinate sum is a fixed translate of $v/T' + j(v)$, where v is uniform modulo T' and $0 \le j(v) \le |S|$. The outer binomial digit remains independent. For each possible integer z the mass of the admissible outer digit is bounded by $\max_{0 \le j \le |S|} w(z-j)$. By telescoping and unimodality,

$$\sum_{z \in \mathbb{Z}} \max_{0 \le j \le |S|} w(z-j) \le 1 + |S| \sum_{j \in \mathbb{Z}} |w(j+1) - w(j)| = 1 + 2|S|\|w\|_\infty.$$

Counting the admissible v for each z gives a conditional local point bound $(1+2|S|\|w\|_\infty)(1/r + 1/T')$. Since $s/\sqrt{\ell} = O_\epsilon(U^{-\epsilon})$, averaging the residual indicators gives $\mathbb{E}(T')^{-1} \ll L^{-3}$ as before. Thus the local estimate (45) holds again. The full numerator range is now $[0, O_\epsilon(D\ell/U)]$, whose wraparound loss has logarithm $O_\epsilon(U)$. The same independence and CRT calculation as in Lemma 6.1 proves (50). □

Theorem 6.6 (Asymmetric contraction of the collision deficit)**.** *For every probability law ζ on $\mathbb{F}_p$, put $d(\zeta) = \log(p\|\zeta\|_2^2) \ge 0$. With ℓ, E_1 as in Lemma 6.5,*

$$d(\zeta * \mu) \le (1 - 1/\ell) d(\zeta) + E_1/\ell. \tag{51}$$

In particular if $d(\zeta) \ge \tau$, adding the anchor once decreases the collision deficit by at least $(1 - o(1))d(\zeta)/\ell$. This is a relative logarithmic decrease at least of order $L/(U^{3+2\epsilon}V)$ at the threshold τ.

Proof. Apply Finite module 4 with $j = \ell$ and $E = E_1$, using Lemma 6.5. This gives (51) with no BSG or comparable-entropy hypothesis. If $d(\zeta) \ge \tau$ then $E_1 = o(d(\zeta))$ uniformly; rearranging gives the stated decrease. □

For a high-entropy heterogeneous example, let H be a rationally dissociated set at height $Q_\sharp = C'LU/\mathcal{C}_4(p)$ supplied by Lemma 2.2, and delete the intersection with A and denominators at most $Q_\sharp/\sqrt{U}$. By the counting upper bound and dissociated-set estimate, choosing C' sufficiently large gives $|H| \ge K'L$ for any prescribed fixed $K' > 0$, and $\sum_{h \in H} 1/h = o(1)$ upon retaining $\lceil K'L \rceil$ elements. Indeed $|A| = o(L)$ and the elementary rational counting bound at $Q_\sharp/\sqrt{U}$ gives $o(L)$ deleted elements. For large p, $Q_\sharp < Q$. Thus $\zeta = \mu_{p,H}$ and μ form a genuine union of distinct denominators with real total cost $o(1)$. Proposition 6.3 and Theorem 6.6 apply to this union, regardless of the modular collisions among the rationally distinct values of H. The higher convolution powers are used only in the proof. No near-full-entropy assertion for H at an arbitrary fixed prime follows from rational dissociation alone.

6.5. **A sharp digit-entropy obstruction at the smaller height.** The use of $Q \asymp LU$ above is essential for this construction. The following bound shows a more general obstruction to obtaining near-full digit capacity at the counting scale using a common denominator of size $p^{O(1)}$. Write $F = \mathcal{C}_4(p)$, so $F \to \infty$, $F = U^{o(1)}$, and $\log F \to \infty$. Fix $C_*, K > 0$, put $Q_\sharp = \lfloor C_* LU/F \rfloor$, and for $J \subseteq [1, Q_\sharp] \cap \mathbb{Z}$ and integer $k' \geq 1$ write

$$\mathcal{V}_{k'}(J) = \left\{ \sum_{d \in J} \frac{c_d}{d} : c_d \in \{0, \dots, k'\} \right\}, \qquad \mathfrak{H}_{K,k'}(p) = \max_{\operatorname{lcm}(J) \leq p^K} \log |\mathcal{V}_{k'}(J)|.$$

The empty set is allowed with lcm one.

Theorem 6.7 (Bounded-common-denominator digit entropy)**.** *For all integers $k' \geq 1$,*

$$\mathfrak{H}_{K,k'}(p) \ll_{C_*,K} \frac{L(1+\log F)}{F} + \frac{L}{U} \log(k'+1). \tag{52}$$

For any fixed $a_* > 1$, uniformly for $\lceil U \rceil \leq k' \leq U^{a_*}$, this is sharp up to constants:

$$\mathfrak{H}_{K,k'}(p) \asymp_{C_*,K,a_*} \frac{L \log F}{F}. \tag{53}$$

The lower bound is achieved by distinct composite denominators whose indicated digit assignments give distinct residues modulo p as well as distinct rational numbers. In particular for every $k' = e^{o(U)}$ a common denominator at most p^K cannot give near-full digit-box capacity at height $Q_\sharp$.

Proof. Put $y = L/F$, which exceeds $\sqrt{Q_\sharp}$ for large p. Split J into denominators with all prime factors at most y and the blocks rm with prime $r > y$, $m \leq Q_\sharp/r$. The lcm of the smooth part has logarithm $O(y + \sqrt{Q_\sharp} \log Q_\sharp) = O(L/F)$. Counting integer numerators in its full digit range shows that its log-cardinality contribution is $O(L/F + \log(k'+1) + \log(2+U))$. The other blocks contain at most $KL/\log y = O_K(L/U)$ different primes, since each divides $\operatorname{lcm}(J) \leq p^K$. For a fixed r, counting with the lcm of $[1, \lfloor Q_\sharp/r \rfloor]$ gives contribution $O(Q_\sharp/r + \log(k'+1))$. For $y < r \leq L$, use

$$\sum_{y < r \leq L,\ r \text{ prime}} r^{-1} \ll (1 + \log F)/U$$

by dyadic prime counts, and for $r > L$ use $Q_\sharp/r \leq C_* U/F$ and the bound on the number of blocks. Summing proves (52); the lower-order terms from the smooth part are absorbed by the displayed right side.

For the lower bound, fix $C_1 > 100(C_* + 1)$ and $c_2 \in (0, \min(K/4, 1/8))$. Take all primes $r \in (C_1 L/F, c_2 L]$, and put $M_r = \lfloor Q_\sharp/r \rfloor$. For each r choose the prime set S_r in $(M_r/2, M_r]$, with product T_r. Uniformly $M_r \to \infty$ and $M_r < U/100$, so $\log T_r \asymp M_r$ and $T_r |S_r| < r$. All denominators rs ($s \in S_r$) are distinct. Their common denominator divides $P_* T_*$ with $P_* = \prod_r r$ and $T_* = \prod_{s \leq U/100,\ s \text{ prime}} s$. By the prime number theorem, $\log(P_* T_*) \leq (c_2 + o(1))L$, hence the common denominator is less than both p^K and $p^{1/4}$ for large p.

Use digits $0 \leq c_{rs} < s$, which are permitted for $k' \geq \lceil U \rceil$. Any collision modulo p has common denominator less than $p^{1/4}$ and real difference at most $U H_{Q_\sharp}$, so its integer numerator has absolute value less than p and must vanish. Reduce the resulting rational relation modulo each outer prime r, after multiplication by $P_* T_*$. As T_*/T_r is invertible modulo r, the local integer $T_r \sum_{s \in S_r} (c_{rs} - c'_{rs})/s$ is zero modulo r. Its absolute value is at most $T_r |S_r| < r$, hence it is zero as an integer. Reducing modulo each s gives equality of the two digits. Thus all $\prod_r T_r$ assignments are distinct modulo p. Their log-cardinality is

$$\sum_r \log T_r \gg \sum_r M_r \gg Q_\sharp \sum_{C_1 L/F < r \leq c_2 L} r^{-1} \gg L \log F/F.$$

The last inequality follows by the prime number theorem on full dyadic subintervals of the displayed prime range. For $k' \leq U^{a_*}$ the second term of (52) is $O_{a_*}(LV/U) = o(L \log F/F)$. This proves (53). Finally the right side of (52) is $o(L)$ for $k' = e^{o(U)}$. Since reduction modulo p cannot increase the number of rational values, the last assertion follows. □

The lower-bound construction also gives a counting-scale asymmetric support-growth statement. Let $A_\sharp$ be its denominator set, with $k_\sharp = \lceil U \rceil$. Its $k_\sharp$-fold support has log-cardinality at least $c_3 L \log F/F$. For any nonempty $Z \subseteq \mathbb{F}_p$, Plünnecke gives

$$\log \frac{|Z + \Sigma_p(A_\sharp)|}{|Z|} \geq \frac{c_3 L \log F/F - \log |Z|}{k_\sharp}.$$

In particular this gives logarithmic gain $\gg L \log F/(FU)$ when $\log |Z| \leq (c_3/2) L \log F/F$. For disjoint denominator sets this is a genuine one-copy gain at height $Q_\sharp$. It is useful only at a sub-full entropy threshold. The obstruction (52) explains why the near-full-capacity proof at height $O(LU)$ cannot be moved to height $O(LU/F)$ by simply replacing the auxiliary block with a different bounded-common-denominator block and retaining $k' = U^{O(1)}$.

Remark 6.8 (Limits of the heterogeneous gains). Theorems 6.4 and 6.6 improve pointwise relative gains and do not establish the absolute target (65). In particular the estimate of Theorem 6.6 can be iterated with independent anchor blocks, but it only guarantees a geometric decay of the deficit at ratio $1 - 1/\ell$. From a deficit of order L, the displayed estimate would require order $\ell \log(L/\tau)$ iterations to guarantee a deficit of order τ; we do not construct that many disjoint anchor blocks at height $O(LU)$. Moreover the sparse block alone has only $2^{O(L/V)}$ binary inputs. Neither its near-full expanded digit capacity nor the support of the rationally dissociated block H implies all-primes binary coverage. The smaller-scale entropy theorem is a restriction on a proof strategy with common denominator at most p^K, not a lower bound ruling out coverage at the counting scale with large common denominator. The new arguments have been checked algebraically within this revision, not by independent third-party referees; no priority claim is certified.

7. Density-one mixing on the auxiliary scale and the packing budget

This section improves the auxiliary convolution order in Section 6. The averaging variable is now a smaller auxiliary prime, not the final modulus. For each fixed $0 < \epsilon < 1$ we obtain a block at height $O_\epsilon(LU)$ whose $O(U^\epsilon)$-fold convolution has maximum atom $p^{-1} \exp(o(L/(UV)))$. The block is available at every sufficiently large final prime. We then establish a packing obstruction for disjoint digital blocks, including blocks with arbitrary common denominators. These results do not give absolute all-primes binary coverage.

7.1. **A local lemma at denominator height proportional to the logarithm.**

For a finite set D of denominators prime to a prime r, write $\mu_{r,D}$ for the law of the independent Bernoulli sum of their inverses in $\mathbb{F}_r$. Use the Fourier transform $\widehat{\mu}(h) = \sum_z \mu(z) \exp(2\pi i h z/r)$, with unnormalized frequency sum.

Lemma 7.1 (Local weak rigidity at logarithmic height)**.** *Let X be large, $x = \log X$ and $u = \log x$. There is an absolute fixed integer M such that, with*

$$R_0 = \{s \text{ prime} : 4x < s \leq 8x\}, \qquad A_0 = \{sm : s \in R_0,\ 1 \leq m \leq M\},$$

we have

$$\#\{(r,h) : X < r \leq 2X,\ r \text{ prime},\ 1 \leq h < r,\ E_0(h,r) \leq 1/u\} \ll X \exp(-cx/u),$$

$$E_0(h,r) = \sum_{a \in A_0} \left\| ha^{-1}/r \right\|^2,$$

for an absolute constant $c > 0$. Here and below the norm denotes distance to the nearest integer.

Proof. For all sufficiently large X, distinct pairs (s,m) give distinct denominators, and all denominators are below X. Put $T_0 = \operatorname{lcm}(1,\dots,M)$ and $n_0 = |R_0| \sim 4x/u$. For a frequency with energy at most $1/u$, let v_s be the centered representative of hs^{-1} modulo r, and write $sv_s = h + rk_s$ with $k_s \in \mathbb{Z}$. Since the denominator s belongs to the block, all but at most $n_0/16$ of the rows have

$$|v_s|/r \leq 4(un_0)^{-1/2}, \qquad |k_s| \leq C_1 \sqrt{x}$$

for an absolute C_1. Call a row internally quiet if for every $1 \le m \le M$ its phase at sm has distance at most $1/(4M)$. At most $16M^2/u = o(n_0)$ rows fail this condition. For a quiet row let v_{sm} be the centered representative of $h(sm)^{-1}$. Then $mv_{sm} \equiv v_s \pmod r$ and $|mv_{sm}| \le r/4$. Centered uniqueness gives $mv_{sm} = v_s$. Consequently $T_0 \mid v_s$, so

$$k_s \equiv -hr^{-1} \pmod{T_0}.$$

At least $7n_0/8$ rows have both properties. The residue $\kappa = -hr^{-1}$ modulo T_0 is common to those rows; for fixed κ and sufficiently large X each such row has at most $C_2\sqrt{x}/T_0$ admissible integers k_s, with C_2 absolute.

Choose $t = \lceil \log(4X^2)/\log(4x) \rceil + 1$ such rows; this is possible since $t \sim 2x/u$. Their prime product exceeds $4X^2$. Each prescribed k_s determines $-hr^{-1}$ modulo s. If another reduced frequency h'/r' gives the same prescribed data, then their prime product divides $hr' - h'r$, which has absolute value less than $4X^2$. Thus the frequencies coincide. The count is at most

$$T_0 2^{n_0} (C_2\sqrt{x}/T_0)^t.$$

For sufficiently large X, $n_0 \le 5x/u$, $t \ge x/u$, and $tu/2 \le x + O(u)$. First choose the fixed M so large that $\log(T_0/C_2) > 10$. Then the logarithm of this count is at most $x - cx/u$. All fixed constants are chosen before X tends to infinity. □

Theorem 7.2 (Local density-one binomial mixing)**.** *Fix $0 < \epsilon < 1$. There are constants $C_\epsilon, c_\epsilon > 0$ such that, for every sufficiently large X, there is a set of denominators independent of r*

$$D_X \subseteq [4x, C_\epsilon x] \cap \mathbb{Z}, \qquad |D_X| \asymp_\epsilon x/u, \qquad \sum_{a \in D_X} 1/a \ll_\epsilon 1/u,$$

where $x = \log X$ and $u = \log x$, and an even integer $j = 2\lceil x^\epsilon/2 \rceil$, with the following property. For all primes $X < r \le 2X$ except at most $O_\epsilon(X \exp(-c_\epsilon x^\epsilon/u))$ exceptions,

$$\mu_{r,D_X}^{*j}(z) = \frac{1}{r}\left(1 + O_\epsilon\left(\exp(-c_\epsilon x^\epsilon/u)\right)\right) \quad (z \in \mathbb{F}_r).$$

Proof. Keep A_0 from Lemma 7.1, and choose a prime denominator block

$$B_0 = \{b \text{ prime} : C_3 x < b \le C_4 x\}$$

with fixed $C_4 > C_3 > 8M$, and $C_4 - C_3 > 10/\epsilon$. Set $D_X = A_0 \cup B_0$; the sets are disjoint. The claimed cardinality and reciprocal-mass bounds follow from the prime number theorem and the fixed parameter choices.

Write $d = j/2$. Let c_b be the difference of two independent $\mathrm{Bin}(d, 1/2)$ variables. The variables c_b are independent over $b \in B_0$, with $|c_b| \le d < b$, and

$$\mathbb{P}(c_b = 0) = 4^{-d}\binom{2d}{d} \ll j^{-1/2}.$$

For $P_0 = \prod_{b \in B_0} b$, put $N = P_0 \sum_{b \in B_0} c_b/b$. Reduction modulo each b shows that $N = 0$ if and only if every $c_b = 0$. Thus, for $n_B = |B_0| = (C_4 - C_3 + o(1))x/u$,

$$\mathbb{P}(N = 0) \le (Cj^{-1/2})^{n_B} \le X^{-3}$$

for large X. For every nonzero N, $\log|N| \ll_\epsilon x$ since $\log P_0 \ll_\epsilon x$ and $j = O(x^\epsilon)$. Such an N has at most $O_\epsilon(1)$ distinct prime factors in $(X, 2X]$. Orthogonality and nonnegativity give

$$\sum_{\substack{X < r \le 2X \\ r \text{ prime}}} \sum_{h \bmod r} |\widehat{\mu}_{r,B_0}(h)|^j = \sum_r r\mathbb{P}(r \mid N)$$

$$\ll_\epsilon X^2\mathbb{P}(N = 0) + X \ll_\epsilon X.$$

For a prime with no nonzero frequency of energy at most $1/u$, the inequality $|\cos(\pi v)| \le \exp(-2\|v\|^2)$ yields $|\widehat{\mu}_{r,A_0}(h)|^j \le e^{-2j/u}$ for every $h \ne 0$. Consequently

$$\sum_{r \text{ good}} \sum_{h \ne 0} |\widehat{\mu}_{r,D_X}(h)|^j \ll_\epsilon Xe^{-2j/u}.$$

By Markov, outside $O_\epsilon(Xe^{-j/u})$ further primes the inner sum is at most $e^{-j/u}$. The initial rigidity exceptions are $O(Xe^{-cx/u})$ and can be absorbed because $\epsilon < 1$. Fourier inversion gives the stated local law. In particular the exceptions are $o(X/\log X)$, not merely $o(X)$. □

7.2. **Deterministic selection of the good auxiliary primes.** Return to $L = \log p$, $U = \log L$, $V = \log U$, and $\tau = c_0 L/(UV)$ for any fixed $c_0 > 0$. Apply Theorem 7.2 at $X = 2L$. It produces a common set of multipliers $D_X \subseteq [4\log(2L), C_\epsilon \log(2L)]$ and a good prime set in $(2L, 4L]$. The exceptional count is $o(L/U)$ and hence the sum of logarithms of the good primes is $(2 + o(1))L$. Select a subset R of good primes by adjoining successive primes until just before their product reaches p. Then

$$p/(4L) < P_R := \prod_{r\in R} r < p, \qquad |R| \asymp L/U.$$

All choices depend on the final p; no exceptional final primes are needed.

Theorem 7.3 (Pointwise auxiliary mixing of arbitrarily small fixed-power order)**.** *For each fixed $0 < \epsilon < 1$ and every sufficiently large prime p, there is a set of distinct denominators*

$$A = A_p = \{ra : r \in R,\ a \in D_X\} \subseteq [1, C'_\epsilon LU], \quad |A| \asymp_\epsilon L/V, \quad \sum_{b\in A} 1/b \ll_\epsilon 1/(UV).$$

Let $\mu = \mu_{p,A}$ and $j = 2\lceil (\log(2L))^\epsilon/2\rceil \asymp U^\epsilon$. There is a nonnegative error bound

$$E_\epsilon(p) = O_\epsilon\left(U + \frac{L}{U}\exp\left(-c_\epsilon \frac{U^\epsilon}{V}\right)\right) = o(\tau)$$

such that

$$\|\mu^{*j}\|_\infty \le p^{-1}\exp(E_\epsilon(p)).$$

Moreover $\operatorname{lcm}(A) \le p\exp(O_\epsilon(U))$.

Proof. The largest prime factor of every ra is r, since $a = O_\epsilon(U) = o(L)$. Thus all denominators are distinct. The size and reciprocal cost follow from the preceding selection and Theorem 7.2. Write $T = \operatorname{lcm}(D_X)$; $\log T = O_\epsilon(U)$ by the standard estimate for $\log\operatorname{lcm}(1,\dots,M)$, and $\gcd(P_R, T) = 1$.

Take independent $\xi_{ra} \sim \mathrm{Bin}(j, 1/2)$, and let $W_r = \sum_{a\in D_X} \xi_{ra}/a$. Theorem 7.2 gives, uniformly in $v \in \mathbb{F}_r$,

$$\mathbb{P}(W_r \equiv v \pmod r) \le r^{-1}(1+\delta), \qquad \delta = O_\epsilon(e^{-c_\epsilon U^\epsilon/V}).$$

The common integer numerator of the global random rational sum is

$$N = P_R T \sum_{r,a} \frac{\xi_{ra}}{ra} = \sum_{r\in R} \frac{P_R}{r} T W_r.$$

Fixing N fixes every W_r modulo r by reduction at r. By independence of rows, $\sup_N \mathbb{P}(N) \le P_R^{-1}(1+\delta)^{|R|}$. Also $0 \le N \le P_R T j \sum_{b\in A} 1/b$. Since p is prime to $P_R T$, reduction of the rational sum is a bijective dilation of $N \bmod p$. Each fiber contains at most $1 + P_R T j \sum_{b\in A} 1/(pb)$ possible integer numerators. The logarithm of this quantity is $O_\epsilon(U)$, and $\log(p/P_R) = O(U)$. The atom bound follows. Finally $\operatorname{lcm}(A) \mid P_R T$, and the stated error is $o(\tau)$ because $U = o(\tau)$ and $V\exp(-c_\epsilon U^\epsilon/V) = o(1)$. □

7.3. **An improved one-copy heterogeneous contraction.** For a probability law ζ on $\mathbb{F}_p$ define $d(\zeta) = \log(p\|\zeta\|_2^2)$. For a nonempty set S put $d_S = \log(p/|S|)$, and write $\Sigma_p(A)$ for the support of $\mu_{p,A}$.

Corollary 7.4 (Asymmetric support and collision-deficit contraction)**.** *With the block and order of Theorem 7.3, for every nonempty $S \subseteq \mathbb{F}_p$ and arbitrary probability law ζ,*

$$\log\frac{|S + \Sigma_p(A)|}{|S|} \ge \max\left(0, \frac{d_S - E_\epsilon(p)}{j}\right),$$
$$d(\zeta * \mu) \le \min\left(d(\zeta),\ (1 - 1/j)d(\zeta) + E_\epsilon(p)/j\right).$$

Thus whenever the input deficit is at least τ, the guaranteed logarithmic gain is $\gg_\epsilon L/(U^{1+\epsilon}V)$. The noise may in particular be a Bernoulli law on any disjoint set of denominators.

Proof. The support of μ^{*j} is exactly $j\Sigma_p(A)$ and its cardinality is at least $pe^{-E_\epsilon(p)}$. If $K = |S+\Sigma_p(A)|/|S|$, Plünnecke's inequality (in its nonempty-subset form) gives $|j\Sigma_p(A)| \le K^j|S|$. This proves the support statement. For the collision assertion, apply Finite module 4 with $E = E_\epsilon(p)$. The corresponding set monotonicity follows by a translation. If the deficit is at least τ, the error is $o(\tau)$, which proves the gain. □

This improves the order $U^{2+2\epsilon}$ in Theorem 6.6 to U^ϵ at a height with an ϵ-dependent constant. In particular it is not legitimate to let ϵ tend to zero with p while keeping this height constant. This is a stronger relative estimate, not a new absolute near-full-entropy bound for the binary law. Taking point-mass input yields an absolute bound of order L/U^ϵ for the collision entropy, which need not improve previously obtained binary entropy bounds.

7.4. **A budget obstruction to packing near-capacity blocks.** The following count explains why one cannot simply iterate the previous corollary using many disjoint copies. For a set J and an integer $k \ge 1$ let $\mathcal{V}_k(J)$ be the rational sumset with digits $0,\dots,k$, as in Section 6.5; its modular image is $k\Sigma_p(J)$ for $J \subseteq [1,Q]$, $Q < p$.

Theorem 7.5 (Packing budget independent of common denominator)**.** *Fix $C, D > 0$ and $Q \le CLU$. Let $J_1,\dots,J_s$ be pairwise disjoint subsets of $[1,Q]\cap\mathbb{Z}$, and let $1 \le k_i \le U^D$. Assume $|k_i\Sigma_p(J_i)| \ge p^{3/4}$ for each i. Then*

$$\sum_{i=1}^{s} \frac{1}{\log(k_i+1)} \ll_{C,D} V, \qquad s \ll_{C,D} V^2.$$

No upper bound on any $\operatorname{lcm}(J_i)$ is assumed.

Proof. Choose a small fixed $\eta > 0$ and put $y = \eta L > \sqrt{Q}$. Partition each J_i into J_i^{small}, with all prime factors at most y, and the complementary J_i^{large}. The least common multiple of all possible small-part denominators has logarithm at most

$$\vartheta(y) + O(\sqrt{Q}\log Q) = (\eta + o(1))L.$$

Their k_i-digit sums lie in $[0, k_iH_Q]$ and on a rational grid with this common denominator. Thus their logarithmic count is at most $L/4$ for sufficiently small η and large p, since $\log k_i = O_D(V)$. All remaining digits contribute at most $|J_i^{\mathrm{large}}|\log(k_i+1)$ to the logarithmic count. Reduction modulo p cannot increase that count, hence

$$|J_i^{\mathrm{large}}|\log(k_i+1) \ge L/2.$$

The total number of integers in $[1,Q]$ with a prime factor exceeding y is at most

$$Q\sum_{y<r\le Q} 1/r \ll_C LV.$$

Indeed this prime reciprocal sum is $O_C(V/U)$ by the prime number theorem with its classical error term and partial summation. If $Q < y$ there is no such denominator and no qualifying block. Disjointness therefore yields the sum bound. The count bound follows from $\log(k_i+1) \ll_D V$. □

We next show that, at the fixed height CLU, the small-power auxiliary order cannot be replaced by a subpower order if one insists on a bounded common denominator. The constants in this assertion depend on the allowed height and common denominator.

Lemma 7.6 (Local subpower digit entropy)**.** *For fixed $C' > 0$ and $M = \lfloor C'U\rfloor$, if $\log(k+1) = o(V)$ then*

$$\log|\mathcal{V}_k([1,M]\cap\mathbb{Z})| = o(U).$$

More quantitatively, for any fixed $0 < \theta < 1$, the left side is at most

$$O(\theta M + \sqrt{M}\log M + \log(k+1) + \log\log(2M)) + O_\theta\left(\frac{M\log(k+1)}{\log M}\right).$$

Proof. Split denominators at the prime factor cutoff θM. The small-prime part has a common denominator of logarithm $O(\theta M + \sqrt{M}\log M)$ and values in an interval of length kH_M. The remaining denominator count is at most $M\sum_{\theta M<r\le M} 1/r = O_\theta(M/\log M)$. Multiply their trivial digit counts. For $\log(k+1) = o(V)$, first fix θ arbitrarily small and let p tend to infinity, and then let θ decrease. This proves the subpower claim. □

Theorem 7.7 (A positive-power order is necessary at bounded common denominator)**.** *For fixed $C, K > 0$ there is a constant $c(C,K) > 0$ such that whenever $J \subseteq [1, CLU]$, $\operatorname{lcm}(J) \le p^K$, and $|k\Sigma_p(J)| \ge p^{3/4}$, one has $k \ge U^{c(C,K)}$ for all sufficiently large p. In particular, if $\log(k+1) = o(V)$ then the logarithm of the rational digit count is $o(L)$, uniformly under these denominator bounds.*

Proof. Fix $\eta > 0$ and put $y = \eta L$. As above, the small-prime part contributes at most $(\eta + o(1))L + O(\log(k+1) + V)$ to the logarithmic count. Every large prime factor $r > y$ occurs to the first power in any denominator, and there are at most $KL/\log y = O_K(L/U)$ distinct such factors in J. The r-block consists of denominators ra with $a \le CLU/r \le (C/\eta)U$, so its rational digit count is bounded by that of $[1, \lfloor (C/\eta)U \rfloor]$ (the factor $1/r$ does not affect counts). Apply Lemma 7.6. If $\log(k+1) = o(V)$, the large-prime blocks contribute $o_{\eta,C,K}(L)$; letting η decrease proves the last assertion. For the positive-power claim choose η so small that the small-prime contribution is less than $L/8$ under $k \le U^c$ with $c \le 1$. Then choose the fixed θ in the quantitative lemma small enough that the $O(\theta M)$ contributions, summed over the $O_K(L/U)$ large primes, are less than $L/8$. Finally choose $c > 0$ small enough that the remaining $O_\theta(M\log(k+1)/\log M)$ contributions are less than $L/8$. Terms of smaller order are negligible. This gives a rational digit count of logarithm less than $L/2$ if $k \le U^c$, contradicting the required modular cardinality. □

Combining Theorems 7.5 and 7.7, any disjoint family of polynomial-in-U order near-capacity blocks at fixed height CLU and fixed common-denominator exponent K satisfies

$$\sum_{i=1}^{s} \frac{1}{k_i} \ll_{C,D,K} V^2 U^{-c(C,K)} = o(1).$$

This is a limitation of the guarantee obtained by successively applying Corollary 7.4 with nonnegative error terms: the product of the certified contraction factors $1 - 1/k_i$ tends to one. It is not a lower bound on the true collision probability of a union. The packing estimate itself permits arbitrary common denominators; the positive-power order lower bound specifically uses the bounded-common-denominator restriction.

7.5. **Where a larger common denominator must enter.** There is a related obstruction to using only separate local maximum-atom estimates, even if the outer prime product is permitted to be much larger than p. Suppose J is composed of rows ra with distinct outer primes $r > \sqrt{Q}$ and $a \le Q/r$, all at height $Q \le CLU$. Use independent digits supported on $\{0, \dots, k\}$, where $\log(k+1) = o(V)$. Denote by ν_r the row law of the inner sum modulo r and set $b_r = \log(r\|\nu_r\|_\infty) \ge 0$. Fix a small $\eta > 0$. For $r > \eta L$, Lemma 7.6 and the support bound imply

$$\log(1/\|\nu_r\|_\infty) \le o_{\eta,C}(U), \qquad b_r \ge (1 - o_{\eta,C}(1))\log r,$$

uniformly in such rows (as $\log r \asymp U$). If the product P_R of all outer primes has $\log P_R \ge L - o(L)$, the total logarithm from rows $r \le \eta L$ is at most $(\eta + o(1))L$. Hence

$$\sum_r b_r \ge (1 - o_{\eta,C}(1))(\log P_R - (\eta + o(1))L).$$

In particular a bound using just $\sup_N \mathbb{P}(N) \le \prod_r \|\nu_r\|_\infty = P_R^{-1}\exp(\sum_r b_r)$ followed by a generic integer wraparound factor cannot certify an error exponent $o(L)$ via that sum of local

deficits at this auxiliary order. The displayed inequality bounds the deficit in this particular estimate, not the actual modular concentration; the row events can have additional correlations when integer translates by p are considered. A stronger version holds for very large outer products even with polynomial-in-U digits. Suppose $k \le U^D$ for a fixed D and $\log P_R/L \to \infty$. For any fixed $K_0 > 1$ and every row $r > K_0 L$, the inner denominator height is $M_r \le CU/K_0$. Its possible rational digit values lie in an interval of length kH_{M_r} on the grid with denominator $\operatorname{lcm}(1,\dots,M_r)$. Thus

$$\log|\operatorname{supp}\nu_r| \le O(CU/K_0) + O_D(V), \qquad b_r \ge \log r - O(CU/K_0) - O_D(V).$$

The total prime logarithm for $r \le K_0 L$ is $O(K_0 L) = o(\log P_R)$ for each fixed K_0. Since $\log r \ge U + \log K_0$ on the other rows, we infer

$$\liminf_{p\to\infty} \frac{\sum_r b_r}{\log P_R} \ge 1 - O(C/K_0).$$

Letting K_0 tend to infinity proves $\sum_r b_r \ge (1-o(1))\log P_R$. Again this is a limitation of the separate-row maximum-atom argument, not of the actual modular distribution. It does not apply to a method which obtains a stronger estimate for the aggregate mass of all integer translates of a residue. A method exploiting large common denominators must estimate the probability carried by those translates more sharply, rather than pay for their number uniformly. This section leaves the original all-primes binary bound and the sharp counting-scale all-primes conjecture open.

8. Growing denominator budgets and multiscale folding fibres

This section addresses the large-common-denominator regime left open in Section 7. It proves a sharp tradeoff for *rational* entropy at growing denominator budgets, and a pointwise packing estimate for the actual joint folding fibres, with heterogeneous row probabilities and short windows. The latter also has a matching obstruction at the level of the weighted-distance information alone. These results do not establish nearly full modular entropy.

8.1. **Rational entropy with a growing denominator budget.** Use L, U, V, W, F_*, G_* and the height range (3) of Section 3. Fix $0 < \gamma < 1$, and allow the real parameter K to depend on p, subject to

$$1 \le K \le \frac{Q}{L\exp(V^\gamma)}. \tag{54}$$

Define

$$\mathcal{E}_K = \frac{QG_*}{U}\log\left(\frac{\log(QV/L)}{\log(Q/(KL))}\right). \tag{55}$$

The argument of the outer logarithm exceeds one. For $K = 1$ this has order $QF_*/(UV)$, agreeing with Theorem 3.1. All implicit constants in this section may depend on the fixed height constants and on γ, but are uniform in the growing parameter K.

Theorem 8.1 (Growing-denominator entropy tradeoff)**.** *Uniformly in (54), with $\mathcal{R}_K(Q,p)$ and $\mathcal{H}_K(Q,p)$ defined as in Section 3,*

$$\log|\mathcal{R}_K(Q,p)| \asymp \mathcal{H}_K(Q,p) \asymp \mathcal{E}_K.$$

For each fixed $0 < \eta < 1$, the lower bound is realized by a dissociated set $H \subseteq [\eta Q, Q]$ with $|H| \asymp \mathcal{E}_K$, $R(H) = o(1)$ and $\log\operatorname{lcm}(H) \asymp KL$ with $\operatorname{lcm}(H) \le p^K$. It admits an outer-prime row decomposition in which the local binary assignments are distinct modulo their respective outer primes. No assertion of modular dissociation at p is made when K grows.

Upper bound. Put $y = L/V$, $D_Q = \operatorname{lcm}(1,\dots,Q)$, $B_0 = \log(Q/(KL))$ and $A_0 = \log(Q/L)$. Thus $A_0 \sim V$, $V^\gamma \le B_0 \le A_0$. The factorization and CRT support argument in the proof of Theorem 3.1 applies without change to growing K. The small-prime factor has logarithm $O(L/V)$. A rational value of reduced denominator at most p^K can have nonzero coordinates

at at most $s = \lfloor KL/\log y \rfloor$ outer primes. The number n_0 of available outer primes is $\asymp Q/U$, so $s \ll n_0 e^{-B_0}$, and the logarithm of the number of choices of their support is

$$O\big(s\log(en_0/s)\big) \ll (Q/U)(1+B_0)e^{-B_0} = o(L/V).$$

This last estimate is uniform since $B_0 \geq V^\gamma$ and $Q \ll LU$. For all real M between $Q/(KL)$ and Q/y, $\log_2 M$ lies between fixed positive multiples of W. Hence $\mathcal{C}_3(M) = \mathcal{C}_2(\log M) \asymp \mathcal{C}_3(U) = G_*$ by (1), or directly by comparing the iterates starting at $\log_3 M = \log(\log_2 M)$. Lemma 2.1 gives $E(M) = \log|\mathcal{V}(\lfloor M \rfloor)| \ll MG_*/\log M$. Include every outer prime between y and KL, and bound the possible remaining selected coordinates by $sE(Q/(KL))$ using monotonicity of E. This yields

$$\log|\mathcal{R}_K(Q,p)| \ll L/V + \log(H_Q+2) + QG_* \sum_{\substack{y<r\leq KL \\ r \text{ prime}}} \frac{1}{r\log(Q/r)} + \frac{QG_*}{UB_0}.$$

The prime sum is $\ll U^{-1}\log((A_0+W)/B_0) + O((UB_0)^{-1})$. For instance partition the interval into fixed-ratio geometric intervals and apply the uniform PNT upper bound $O(t/U)$ on each interval of scale t; the corresponding integral is $U^{-1}\int_y^{KL} dt/(t\log(Q/t))$. All $\log t \asymp U$. The endpoint error $QG_*/(UB_0)$ is absorbed in $\mathcal{E}_K$: indeed $B_0\log((A_0+W)/B_0)$ tends uniformly to infinity (if $B_0 \geq A_0/2$ use $W \to \infty$, otherwise use $B_0 \geq V^\gamma$). Also $\mathcal{E}_K \gg L/V$, by $F_* = WG_*$ and $K \geq 1$. This proves the upper bound. The common-denominator entropy cannot exceed the varying-denominator entropy. □

Lower bound and construction. Choose a small fixed $\alpha > 0$ and a large fixed ratio $A > 1$, and use geometric scales

$$Y_j = \alpha KLA^{-j}, \qquad 0 \leq j \leq J = \left\lfloor \frac{\log(K\sqrt{V})}{\log A} \right\rfloor, \qquad M_j = \lfloor Q/Y_j \rfloor.$$

The inner scales range from constant multiples of $Q/(KL)$ up to $O(U\sqrt{V})$. Consequently $\mathcal{C}_3(M_j) \asymp G_*$ uniformly, and $M_jG_*/\log M_j = o(U)$ uniformly (use monotonicity of $M/\log M$ and $G_* = V^{o(1)}$). As in the proof of Theorem 3.1, the dissociated-set estimate of Lemma 2.2 and Lemma 2.1 give a dissociated set of order $M_jG_*/\log M_j$ in $[c_1M_j, M_j]$ for fixed small $c_1 > 0$. Partition this interval into a bounded number of intervals of ratio $1+\epsilon$, with $(1+\epsilon)^{-2} > \eta$. Retain one subblock $D_j \subseteq [s_j, (1+\epsilon)s_j]$, still of that order. Choose outer primes in

$$[Q/((1+\epsilon)^2 s_j),\ Q/((1+\epsilon)s_j)].$$

These intervals have fixed ratio $1+\epsilon$ and are comparable to Y_j; by taking A large enough they are pairwise disjoint. Their primes exceed all inner denominators. Each contains $\asymp Y_j/U$ primes. The number at which some nontrivial signed relation on D_j vanishes is at most $3^{|D_j|}O(1+M_j/U) = \exp(o(U))$, uniformly: each nonzero rational relation has, after clearing denominators, a nonzero numerator with logarithm of its absolute value $O(M_j)$. This is negligible compared with $Y_j/U \gg L/(U\sqrt{V})$. Retain the good primes, denoted R_j.

Put $R = \bigcup_j R_j$, $H = \bigcup_j \{ra : r \in R_j, a \in D_j\}$, $T = \operatorname{lcm}(1,\dots,\max_j M_j)$ and $P_R = \prod_{r\in R} r$. The denominators are distinct, belong to $[\eta Q, Q]$, and their common denominator divides TP_R. Moreover $(T, P_R) = 1$, $\log T = O(U\sqrt{V}) = o(L)$, and $\log P_R \asymp \alpha KL$ (the first scale gives the lower bound; the PNT applied to all primes below a constant multiple of αKL gives the upper bound). Each outer prime divides the actual least common multiple. Choose α small enough to have $\log(TP_R) \leq KL$ for all sufficiently large p. Local injectivity modulo r implies exact rational dissociation of H: reduce any putative signed equality modulo each outer prime after clearing the common denominator; only that row survives, and all coefficients in the row vanish.

The cardinality is

$$|H| \asymp \frac{QG_*}{U}\sum_{j=0}^{J}\frac{1}{\log M_j} \asymp \frac{QG_*}{U}\log\left(\frac{A_0+W}{B_0}\right).$$

Indeed the inner logarithms increase in steps of $\log A+o(1)$ from $B_0+O(1)$ to $A_0+W/2+O(1)$; integral comparison, with endpoint errors $O(B_0^{-1})$, proves the last comparison uniformly since $W \to \infty$. Replacing $W/2$ by W changes the comparison by at most a fixed factor. Here $R(H) \ll |H|/Q = O(G_*W/U) = o(1)$, since the outer logarithm in (55) is $O(W)$ and $F_* = WG_* = V^{o(1)}$. Thus $2^{|H|}$ distinct rational values realize the lower bound. □

Corollary 8.2 (Where a rational entropy budget of size L first becomes possible)**.** *At height $Q \asymp LU$, define K_{rat} as the infimum of the $K \geq 1$ for which $\mathcal{H}_K(Q,p) \geq L$. Then*

$$\log K_{\mathrm{rat}} \asymp V/G_*.$$

The same estimate holds if $\mathcal{H}_K$ is replaced by $\log|\mathcal{R}_K|$. At height $Q \asymp LU/F_*$, for every fixed $\delta > 0$ one has $\log|\mathcal{R}_K| = o(L)$ when $1 \leq K \leq U^{1-\delta}$. For sufficiently large fixed C in $Q = \lfloor CLU/F_* \rfloor$, there is nevertheless a row-dissociated H with $|H|\log 2 \geq L$ and $\log \operatorname{lcm}(H) \ll Qe^{-\sqrt{V}}$. In particular at that sharp height the least K allowing rational entropy L has $\log K = (1+o(1))V$; this is not a statement about the least K allowing modular coverage.

Proof. For $Q \asymp LU$ take $K = \exp(cV/G_*)$ with a fixed $c > 0$. Then $\log K = o(V)$, $W = o(V/G_*)$, and $\mathcal{E}_K \asymp cL$. Taking c small or large, respectively, proves the threshold by Theorem 8.1 and monotonicity. At the sharp height, $K \leq U^{1-\delta}$ gives $B_0 \gg_\delta V$, hence $\mathcal{E}_K \ll_\delta QG_*/U \ll_\delta L/W = o(L)$. For the construction choose $\gamma = 1/2$ and $K = Q/(Le^{\sqrt{V}})$. Then the outer logarithm is $(1/2+o(1))W$, and the lower construction gives $|H| \gg QF_*/U \asymp CL$. The implied lower constant can be taken independent of the fixed C for sufficiently large p depending on C: the scale comparisons for M_j and A_0 are asymptotically unaffected by $\log C$. Choose C sufficiently large for the stated cardinality. The logarithm of its least common denominator is $O(KL) = O(Qe^{-\sqrt{V}})$. These bounds also imply the result for the varying-denominator version. □

Corollary 8.3 (A reduced-denominator restriction for almost every target at each prime)**.** *For each sufficiently large prime p and $Q \asymp LU$, if $K \geq 1$ and $\log K = o(V/G_*)$, then at most $p^{o(1)}$ residue classes have a reciprocal subset-sum representation in $[1,Q]$ whose reduced denominator is at most p^K. At $Q \asymp LU/F_*$ the same assertion holds for $K \leq U^{1-\delta}$ with any fixed $0 < \delta < 1$. This is a statement about almost all targets at a given prime, not about almost all primes, and not a lower bound for any prescribed target.*

Proof. In the first case $\log K = o(V)$ and the outer logarithm in (55) is $O((W+\log K)/V)$; hence $\mathcal{E}_K = o(L)$. The second case was proved in Corollary 8.2. Rational reduction cannot increase the number of values, since every allowed denominator is prime to p. Thus the exceptional target count is bounded by $|\mathcal{R}_K(Q,p)| = \exp(o(L))$. □

8.2. **Joint folding in a window: a heterogeneous weighted packing bound.** Consider distinct outer primes R below p, a positive integer T coprime to pP_R, and independent random row values X_r with denominators dividing T. In each row suppose the finite support is injective on reduction modulo r. Let w_r be its probability law and $a_r = \|w_r\|_\infty$. Define $S = \sum_r X_r/r$ as a rational random variable, $D = TP_R$, and the integer numerator $N = DS$. For $x \in \mathbb{F}_p$ write n_x for the representative of Dx in $[0, p-1]$. On $S \equiv x \pmod p$ the folding index $z = (N-n_x)/p$ is an integer.

For $s \geq 0$ define the weighted ball expression

$$\mathcal{B}(s) = \sum_{J \subseteq R,\ \sum_{r\in J}\log r \leq s} \prod_{r\in J} (a_r^{-1}-1). \tag{56}$$

This expression retains the heterogeneity of the rows and is independent of the final modulus except through the chosen rows and threshold.

Theorem 8.4 (Joint folding in short windows)**.** *For any $Z \geq 1$ with $Z < P_R$, any real interval I of length Z, and any $x \in \mathbb{F}_p$, one has, for $0 \leq s < \frac{1}{2}\log(P_R/Z)$,*

$$\mathbb{P}(S \equiv x \pmod p,\ z \in I) \leq \mathcal{B}(s)^{-1}.$$

If the range of S is contained in a real interval of length less than one and $T < p$, then for $0 \le s < \frac{1}{2}\log(p/T)$ the global pointwise estimate $\mathbb{P}(S \equiv x \pmod p) \le \mathcal{B}(s)^{-1}$ holds as well. The result allows highly nonuniform row laws and arbitrarily large P_R.

Proof. The map from row values to N is injective: modulo each r, N determines $X_r \pmod r$, hence its actual row value. Let two different row vectors produce the same residue x and folding indices z, z' in I. For every row on which they agree, $r \mid N - N' = p(z - z')$. Since $p \nmid P_R$ and $0 < |z - z'| \le Z$, the product of the agreeing outer primes is at most Z. Thus their weighted Hamming distance, with row cost $\log r$, is at least $\log(P_R/Z)$. Globally, if the changed row set is J, their nonzero rational difference has denominator dividing $T\prod_{r\in J} r$ and absolute value less than one; if that denominator bound is at most p, reduction modulo p cannot be zero. Hence the weighted distance exceeds $\log(p/T)$ in the global case.

The balls of radius s around vectors in either fibre are disjoint by the triangle inequality. If a centre has product weight w, changing precisely the coordinates in J gives total weight $w\prod_{r\in J}(1 - w_r(X_r))/w_r(X_r) \ge w\prod_{r\in J}(a_r^{-1} - 1)$. Therefore each ball has mass at least $w\mathcal{B}(s)$. Their total probability is at most one, proving both assertions. For an interval containing at most one possible folding index, the fibre has at most one row vector and the bound $\prod_r a_r$ also holds. □

Corollary 8.5 (Multiscale folded-mass bounds for the constructed high-entropy blocks)**.** *For the set H from Theorem 8.1, let S be its Bernoulli reciprocal sum, and use its row factorization T, P_R from the construction. For every $1 \le K' \le K$ and every real interval I' of length at most $p^{1-K'}/T$,*

$$\sup_{x\in\mathbb{F}_p} \mathbb{P}(S \equiv x \pmod p,\ S \in I') \le \exp(-c\mathcal{E}_{K'}).$$

The global collision bound is $\rho_p(H) \le \|\mu_{p,H}\|_\infty \le \exp(-c\mathcal{E}_1)$, regardless of the growing common-denominator budget. In particular, the latter exponent has the same order as the earlier bounded-denominator pointwise exponent, not a new nearly full entropy estimate.

Proof. Each row law is uniform on 2^{d_r} values, $d_r = |D_j|$ for $r \in R_j$. Thus $a_r = 2^{-d_r}$. A real interval of the specified length gives a folding interval of length at most $P_R p^{-K'}$. If this is less than one the one-point estimate applies. Otherwise take $s = K'L/3$ in Theorem 8.4.

For all such K', there are complete geometric scales in the construction with Y_j below $c_2K'L$ and down to the smallest scale $\asymp L/\sqrt{V}$, where $c_2 > 0$ is a sufficiently small fixed constant (at most α). Their total outer log product is at most $K'L/3$ by the PNT and geometric spacing. All variations on these rows are permitted in the ball. The calculation of $|H|$ in Theorem 8.1, now with upper scale comparable to $K'L$, gives $\sum d_r \gg \mathcal{E}_{K'}$ on the selected rows. The change of either endpoint by a fixed factor changes the logarithmic integral by $O(1/B_0)$ and is absorbed uniformly; the scale ratio between upper and lower endpoints is at least a constant multiple of $\sqrt{V}$. Hence $\mathcal{B}(K'L/3) \ge \exp(c\mathcal{E}_{K'})$. The global bound follows using a fixed radius, for instance $L/4$, since $\log T = o(L)$ and $R(H) = o(1)$. Collision probability is at most maximum point mass. □

8.3. **What the weighted-distance constraint cannot prove.** For the same constructed block, write $\mathcal{B}_H(s)$ for (56) with $a_r = 2^{-d_r}$. Uniformly over (54),

$$\log \mathcal{B}_H(L) \asymp \mathcal{E}_1 = o(L). \tag{57}$$

The lower bound follows by the small-scale selection in Corollary 8.5. For the upper bound a set of changed rows of weight at most L has at most $O(L/U)$ rows, since every outer prime exceeds a fixed multiple of $L/\sqrt{V}$. The logarithm of the number of such supports is $O(LV/U)$ (there are $O(L)$ available primes). For rows up to L the sum of their d_r is at most

$$\ll \frac{QG_*}{V} \sum_{\substack{L/V<r\le L\\ r \text{ prime}}} \frac{1}{r} \ll \frac{QG_*W}{UV},$$

using the same local entropy estimate as before. Above L each selected row has $d_r \ll E(Q/L) \ll (Q/L)G_*/V$ by monotonicity, and there are $O(L/U)$ of them. Combining proves (57), including the contribution $O(LV/U) = o(\mathcal{E}_1)$.

A distance-only obstruction now uses the *actual local alphabets*, not artificially chosen field coefficients. On the product of the row alphabets connect two distinct vectors whenever their weighted Hamming distance is at most L. Each vertex has degree $\mathcal{B}_H(L) - 1$. Greedy colouring therefore gives a class of product probability at least $1/\mathcal{B}_H(L) = \exp(-O(\mathcal{E}_1))$ whose weighted minimum distance exceeds L. Such a class satisfies all the pairwise distance restrictions used in the global part of Theorem 8.4. This remains true when the total binary input entropy $|H| \log 2$ is much larger than L. This class is not asserted to be a modular fibre; the example proves only that the distance constraints and heterogeneous row entropies alone cannot imply the required bound $p^{-1} \exp(o(L/(UV)))$.

For clarity, the desired joint estimate for an actual row-dissociated H of binary input entropy $d \log 2$ is an upper bound of the form

$$\#\{z \in \mathbb{Z} : n_x + pz \text{ is a numerator arising from } H\} \le 2^d p^{-1} \exp(o(L/(UV)))$$

for each $x \in \mathbb{F}_p$, with the natural feasible range of z. Theorem 8.4 controls shorter windows with a larger denominator-distance threshold, but summing such windows loses their large number; it does not establish this global estimate. Equidistribution theorems for general CRT sets, such as Kowalski–Soundararajan [17], are useful context but average over composite moduli, and are not invoked as a pointwise theorem about the specific twist and exponentially short relative windows here. A successful next step must exploit the simultaneous arithmetic constraints on long changes of rows, beyond the distance criterion. The all-primes covering conjecture remains open in this manuscript.

9. An asymptotically sharp entropy budget at the recursive scale

The preceding construction uses a dissociated subfamily and so establishes only comparability. We now use distinct rational row values, rather than binary input vectors, to obtain an asymptotic equivalence. This refinement is based on the recursion established in the companion manuscript [18, entropy theorem; source label `thm:entropy`]; we recall the necessary argument to make the dependence explicit. No assertion of asymptotic Bernoulli entropy, or of full modular coverage, is implicit in this section.

Write $s(x) = \log |\mathcal{V}(\lfloor x \rfloor)|$ and

$$A(x) = \sum_{1 \le j \le x} \frac{s(j)}{j(j+1)}.$$

The recursion and its order estimates are

$$s(x) = \frac{x}{\log x}(A(\log x) + O(1)), \qquad s(x) \asymp \frac{x}{\log x}\mathcal{C}_3(x). \tag{58}$$

For the first formula, split at the prime threshold $y = x/(c \log x)$ for a small fixed $c > 0$. Each large prime $r > y$ occurs in a disjoint block $r[1, \lfloor x/r \rfloor]$. In an equality of two reciprocal subset sums, reduction after multiplying by r gives the local signed relation modulo r. On clearing the inner denominators this is an integer of magnitude $\le \operatorname{lcm}(1, \dots, \lfloor x/r \rfloor) H_{\lfloor x/r \rfloor} < r$, so it vanishes exactly. Thus the distinct rational values factor into the local distinct-value sets and the core. The core entropy is $O(x/\log x)$ by its common denominator bound. The PNT, uniformly on the intervals $(x/(j+1), x/j]$ for $j \le \log x$, gives $x/(j(j+1)\log x)(1 + O(\log_2 x/\log x))$ primes. Summing, using $s(j) \le j \log 2$, and extending the sum from $c \log x$ to $\log x$ at normalized cost $O_c(1)$, proves the first formula. The upper bound in the second formula is Lemma 2.1, and the lower bound follows from Lemma 2.2. In particular $A(x) \to \infty$, and $s(tx) \sim ts(x)$ for fixed $t > 0$, since changing the argument of A by a fixed additive constant changes it by $O(1/\log x)$ when its argument is $\log x$. With $b(P) = \min\{n : s(n) \ge \log P\}$ this implies $s(\lfloor tb(P) \rfloor) \sim t \log P$ and $b(P) \asymp \log P \log_2 P/\mathcal{C}_4(P)$.

Lemma 9.1 (Uniform local entropy normalization)**.** *For every fixed $0 < \gamma < 1$, with L, U, V, W as above,*

$$s(M) = (1 + o(1))\frac{MA(V)}{\log M}$$

uniformly when $V^\gamma/2 \le \log M \le V + O(W + \log W)$. Moreover $A(V) \asymp G_*$, $A(U) \sim WA(V)$ and, throughout (3),

$$s(Q) \sim \frac{QWA(V)}{U}. \tag{59}$$

Proof. By (58) and the bounded-error recursion, $A(V) \asymp \mathcal{C}_3(U) = G_* \to \infty$. For $a = V^\gamma/2$ and $b = V + O(W + \log W)$, monotonicity and the counting upper bound give

$$0 \le A(b) - A(a) \ll \mathcal{C}_3(2V) \sum_{a<j\le b} \frac{1}{j \log j} = O_\gamma(\mathcal{C}_3(V)).$$

But $\mathcal{C}_3(V) = G_*/\log_2 V = o(G_*)$. This proves the first assertion. To see $A(U) \sim WA(V)$, put $j_0 = \lceil \exp(V^\theta) \rceil$, with fixed $0 < \theta < 1$. On $j_0 \le j \le U$, the first assertion (with exponent θ) and integral comparison give

$$\sum_{j_0\le j\le U} \frac{s(j)}{j(j+1)} = (1 + o(1))(1 - \theta)WA(V).$$

For the upper bound use the recursion uniformly for all j above a fixed cutoff, $A(\log j) \le A(V)$, and $\sum_{2\le j\le U} 1/((j + 1)\log j) = W + O(1)$. The finitely many terms and the $O(1)$ error in (58) contribute $o(WA(V))$. Let θ decrease to zero. Finally $\log Q = U + O(V)$, and $A(\log Q) - A(U) = O(V/U)$ by $s(j) \le j \log 2$. Apply (58) once more to obtain (59). □

9.1. **Finite CRT counting interface for varying reduced denominators.**

Finite module 5 (Outer-prime row counting)**.** *Let $Q \ge 2$ be an integer, $\sqrt{Q} < y < Q$, $D_Q = \operatorname{lcm}(1, \dots, Q)$, $R_y = \{r \text{ prime} : y < r \le Q\}$, $P_y = \prod_{r\in R_y} r$, and $T = D_Q/P_y$. Write $S(j) = |\mathcal{V}(j)|$, $M_r = \lfloor Q/r \rfloor$. For any $v \in \mathcal{V}(Q)$ put $n_v = D_Q v \in \mathbb{Z}$. Then*

$$r \mid \operatorname{den}(v) \iff n_v \not\equiv 0 \pmod r \quad (r \in R_y). \tag{60}$$

If $\operatorname{den}(v) \le B_0$ with real $B_0 \ge 1$, its outer support contains at most $m = \lfloor \log B_0/\log y \rfloor$ primes. Moreover

$$\#\{v \in \mathcal{V}(Q) : \operatorname{den}(v) \le B_0\} \le T(\lfloor H_Q \rfloor + 1) \sum_{\substack{I\subseteq R_y \\ |I|\le m}} \prod_{r\in I} S(M_r). \tag{61}$$

For the lower construction, take any set $R \subseteq R_y$ such that the distinct values of $\mathcal{V}(M_r)$ are injective modulo r for each $r \in R$. The denominator set $J = \bigcup_{r\in R} r[1, M_r]$ has distinct elements and

$$|\{R(J') : J' \subseteq J\}| = \prod_{r\in R} S(M_r), \qquad \operatorname{lcm}(J) \mid \operatorname{lcm}(1, \dots, \max_{r\in R} M_r) \prod_{r\in R} r. \tag{62}$$

The empty case uses empty products and lcm equal to one.

Proof. Since $r^2 > Q$ and the product of two different primes in R_y exceeds Q, their rows are disjoint, each occurs to exponent one in D_Q, and $(T, P_y) = 1$. Reducing n_v/D_Q to lowest terms proves (60). Each prime in the support exceeds y, and their product divides the reduced denominator, which proves the support bound. Fix a representation of v by a subset of $[1, Q]$. Modulo r, the numerator n_v equals (D_Q/r) times the modular value of that row's inner subset sum. All other contributions vanish. Thus its residue has at most $S(M_r)$ possible values, even when two different inner rational values collide modulo r. The inactive outer coordinates are zero. Modulo T there are at most T possibilities. The Chinese remainder theorem then bounds the number of possible residues of n_v modulo D_Q by the right side of (61) without the factor $\lfloor H_Q \rfloor + 1$. For a fixed such residue the values v differ by integers and belong to $[0, H_Q]$, giving that factor. In the lower construction every row denominator divides $T_0 \prod_{r\in R} r$,

where $T_0 = \operatorname{lcm}(1, \dots, \max M_r)$ is prime to all outer primes. Clear denominators by this product in any equality between two tuples of local rational values. Modulo each r it forces equality of that row's residues, and hence its actual local values by hypothesis. This proves injectivity of the tuple-sum map; it is also surjective onto the displayed sumset by definition. The denominator divisibility is immediate. Note that this is a bijection of tuples of *distinct rational values*, not of Bernoulli input vectors. □

9.2. **Uniform asymptotic evaluation of the finite ledger.**

Theorem 9.2 (Sharp growing-denominator entropy profile)**.** *Fix the constants in* (3) *and* $0 < \gamma < 1$. *Uniformly for the real* K *in* (54)*, put*

$$\mathcal{F}_K = \frac{QA(V)}{U} \log\left(\frac{\log(QV/L)}{\log(Q/(KL))}\right).$$

Then

$$\log|\mathcal{R}_K(Q,p)| \sim \mathcal{H}_K(Q,p) \sim \mathcal{F}_K. \tag{63}$$

The lower bound can be realized by all multipliers up to $\lfloor Q/r \rfloor$ *in suitable outer-prime rows, with total real reciprocal mass* $o(1)$ *and common denominator at most* p^K*. The local* distinct rational values *are injective modulo their respective outer primes. For* $K = 1$ *the constructed rational values are also injective modulo the final prime* p*. For growing* K*, no such final-prime injectivity is asserted. In contrast to Theorem 8.1, this construction need not be a dissociated binary set or lie in a fixed upper subinterval of* $[1, Q]$.

Proof. Let $a = \log(Q/L) \sim V$, $b = \log(Q/(KL)) \in [V^\gamma, a]$, and $I = \log((a+W)/b)$. Thus $I \gg W/V$, $\mathcal{F}_K \gg L/V$, and $\log K = O(V)$. All $o(1)$ terms in this proof are uniform in the permitted Q, K. Here the lower bound on $\mathcal{F}_K$ uses $Q \gg LU/F_*$ and $A(V) \asymp G_* = F_*/W$. The uniform quantifier is: fix c_*, C_*, γ and a positive error tolerance first, then choose a single lower bound on p valid for every permitted Q, K. In particular b is never allowed to tend to zero; its lower bound V^γ is needed for the following support and endpoint estimates.

For the upper bound use the cutoff $y = L/(VW)$, which exceeds $\sqrt{Q}$. The small-prime factor of $D_Q = \operatorname{lcm}(1, \dots, Q)$ has log size $O(y + \sqrt{Q}\log Q) = o(L/V)$. Finite module 5 with $B_0 = p^K$ gives a support bound $m = \lfloor KL/\log y \rfloor$ and the exact counting inequality (61). There are $n \asymp Q/U$ available primes, with $m/n = O(e^{-b}) = o(1)$. For $m \le n/2$, $\sum_{j \le m} \binom{n}{j} \le (m+1)(en/m)^m$ when $m \ge 1$ (and equals one if $m = 0$). This follows by monotonicity up to m and the standard binomial bound. The logarithm of the number of possible coordinate supports is

$$O(m \log(en/m)) = O((Q/U)(1+b)e^{-b}) = o(L/V).$$

For each support its log number of possible numerator patterns, apart from the small-prime factor, is at most the sum of $s(Q/r)$ over its primes. The weight $s(Q/r)$ is nonincreasing in r. Therefore this sum is maximized by taking the first m primes greater than y. Let $\delta = V^{-1/4}$. By the PNT, the number of primes in $(y, (1+\delta)KL]$ exceeds m for large p, since $y/(KL) \ll 1/(VW)$ and $\log(KL)/\log y = 1 + O(V/U)$. Thus it suffices to sum over all primes in that interval. The remaining multiplicative factors in (61) contribute $\log T = o(L/V)$ and $\log(\lfloor H_Q \rfloor + 1) = O(V) = o(L/V)$; neither depends on the chosen support. Lemma 9.1 applies to all the relevant inner heights. By partial summation and $\log r \sim U$ uniformly,

$$\begin{aligned}\sum_{y < r \le (1+\delta)KL} s(Q/r) &= (1+o(1))\frac{QA(V)}{U}\log\left(\frac{a+W+\log W}{b-\log(1+\delta)}\right)\\ &= (1+o(1))\mathcal{F}_K.\end{aligned}$$

For the analytic step set $w(t) = 1/(t\log(Q/t))$. On this interval $\log(Q/t) \ge b - O(\delta) > 1$, hence $w'(t) = -w(t)(1 - 1/\log(Q/t))/t$. For $E(t) = \pi(t) - \operatorname{li}(t)$ the classical PNT error gives $|E(t)| \ll te^{-c\sqrt{U}}$ uniformly. Stieltjes partial summation bounds the error in replacing $\sum w(r)$ by $\int w(t)\,dt/\log t$ by endpoint errors plus $\int |E(t)w'(t)|dt$, at most $O(e^{-c\sqrt{U}}(1/b + 1 + I))$. The

integral $\int w(t)dt$ is the exact logarithmic ratio displayed above; also $1/\log t = (1+O(V/U))/U$ uniformly. The error is negligible compared with I/U since $I \gg W/V$. The changes of endpoints in the logarithmic integral are $O(\log W/V + \delta/b) = o(I)$: if $b \geq a/2$ use $I \gg W/V$, and otherwise I is bounded below by a positive constant. This proves the upper bound, including for varying reduced denominators.

For the lower bound let $\epsilon = W^{-1/2}$, $y' = L/V^{1-\epsilon}$, and $z' = (1-\delta)KL$. For $y' < r \leq z'$ put $M_r = \lfloor Q/r \rfloor$. The maximum inner height is $O(UV^{1-\epsilon})$, and Lemma 9.1 gives

$$s(M_r) \ll UG_*V^{-\epsilon} = o(U),$$

using $\log G_* = O(\log W)$ and $V^\epsilon = \exp(\sqrt{W})$. Divide $(y', z']$ into dyadic pieces. In each use the maximum inner height M there. Distinct rational sums from $[1, M]$ give at most $e^{2s(M)}$ ordered nonzero differences. Clearing them by $\operatorname{lcm}(1, \dots, M)$ gives nonzero integers of log magnitude $O(M)$; each has $O(1 + M/U)$ possible prime divisors in the piece. The number of primes at which any two distinct local values coincide is thus $\exp(o(U))$ on that piece, uniformly. There are $O(V)$ pieces, so removing all these primes costs at most $\exp(o(U)) = o(L/V)$ in total local entropy (the factor $\max_r s(M_r) = o(U)$ is absorbed). For every remaining prime r all rational values in $\mathcal{V}(M_r)$ are distinct modulo r.

Let $J = \bigcup_r r[1, M_r]$ over the remaining primes. Its denominators are distinct and at most Q. Its common denominator divides TP_R, where $T = \operatorname{lcm}(1, \dots, \max_r M_r)$ and $P_R = \prod_r r$. The log of T is $O(UV^{1-\epsilon}) = o(\delta L)$ and $\log P_R \leq \vartheta(z') = (1 - \delta + o(\delta))KL$, so $TP_R < p^K$. Its reciprocal mass is bounded by

$$\sum_{y' < r \leq z'} H_{M_r}/r \ll V \frac{\log(z'/y')}{U} = O(V^2/U) = o(1).$$

By Finite module 5, equation (62), the rational sumset of J has exactly the product of the local cardinalities. No assumption of Bernoulli injectivity is required. The same partial summation as above gives its log size

$$(1 + o(1))\frac{QA(V)}{U} \log\left(\frac{a + (1-\epsilon)W}{b - \log(1-\delta)}\right) = (1 + o(1))\mathcal{F}_K.$$

The removed prime contribution is negligible. The endpoint changes are $O(\epsilon W/V + \delta/b) = o(I)$ as before. This proves the matching lower bound for the common-denominator entropy. The inequalities between the two entropy definitions then imply (63). If $K = 1$, $TP_R < p$ and $R(J) < 1$, so no nonzero rational difference can vanish modulo p; its integer numerator would be nonzero with absolute value strictly less than p. This proves the last injectivity assertion. □

Corollary 9.3 (Exact budget threshold above the entropy scale)**.** *Let $t > 1$ be fixed and $Q = \lfloor tb(p) \rfloor$. For every fixed $0 < \beta < 1$, put $K_\beta = Q/(L\exp(V^\beta))$. Then*

$$\log|\mathcal{R}_{K_\beta}(Q, p)| \sim \mathcal{H}_{K_\beta}(Q, p) \sim t(1-\beta)L.$$

In particular let K_{rat} be the least real $K \geq 1$ for which $\mathcal{H}_K(Q, p) \geq L$. Then

$$\log\left(\frac{Q}{K_{\mathrm{rat}}L}\right) = V^{1-1/t+o(1)}. \tag{64}$$

The same threshold holds for the least K with $\log|\mathcal{R}_K(Q, p)| \geq L$. For fixed $\beta > 1 - 1/t$, the fraction of residues reachable with reduced denominator at most p^{K_β} is at most $p^{t(1-\beta)-1+o(1)}$ at every sufficiently large prime.

Proof. Use Theorem 9.2 with any fixed $\gamma < \beta$. Its logarithmic integral is $\log((a+W)/V^\beta) = (1-\beta)W + o(1)$ since $a \sim V$. By (59) and regular variation, $QWA(V)/U \sim s(Q) \sim tL$. This gives the two asymptotics. For any sufficiently small fixed $\varepsilon > 0$, take $\beta_\pm = 1 - 1/t \pm \varepsilon$ in $(0, 1)$. The budget K_{β_-} has entropy strictly exceeding L for large p, whereas K_{β_+} does not. Monotonicity gives $K_{\beta_+} < K_{\mathrm{rat}} \leq K_{\beta_-}$, and the same for varying denominators. Taking logarithms and

letting $\varepsilon \downarrow 0$ yields (64). The target bound follows because reduction cannot increase cardinality. The thresholds exist (finite sets, and the displayed larger budget suffices). □

Corollary 9.4 (Normalization at the larger height)**.** *If $Q \sim CLU$ with fixed $C > 0$, then the least K for either rational entropy to reach L satisfies*

$$\log K \sim \frac{V}{CA(V)}.$$

Proof. For fixed $c > 0$ take $\log K = cV/(CA(V))$. Since $A(V) \asymp G_* = V^{o(1)}$ tends to infinity, this is $o(V)$ but much larger than W. The parameter is in (54), and $\log((a+W)/(a-\log K)) \sim \log K/V$. Hence the entropy in Theorem 9.2 is $(c + o(1))L$. Sandwich between $c = 1 - \varepsilon$ and $c = 1 + \varepsilon$. □

The last two corollaries sharpen the rational capacity thresholds in Corollary 8.2, not the actual modular coverage height. When the allowed common denominator exceeds p, many distinct rational values can fold into the same residue. Even at entropy strictly above L no upper bound near p^{-1} for the maximum modular atom follows from the cardinality. The distinct-value construction here and the dissociated construction of Theorem 8.1 therefore serve different purposes. The former yields a leading asymptotic constant; the latter supplies independent Bernoulli row laws for the weighted folding estimates.

10. The pointwise covering problem and limitations

The all-primes coverage and arbitrary-denominator height results of [18] retain a factor $\log_3 p/\mathcal{C}_5(p)$ over the scale $\log p \log_2 p$. The matching progression theorem in that work shows that a single rational grid of length at least p cannot improve this order. This is an obstruction to a proof strategy, not a lower bound excluding $g(p) \ll \log p \log_2 p$ itself. At the smaller counting scale the same distinction is necessary: the exact rational budget of Theorem 9.2 is a necessary information budget, not a sufficient modular-coverage condition. Modular coverage alone also does not yield an exact rational height theorem for arbitrary denominators.

Return to the dyadic blocks of Section 2.1, and take B to be the prime block used in Proposition 3.3. A concrete sufficient pointwise target is

$$\max_{\substack{X<p\le 2X \\ p \text{ prime}}} \log(p\rho_p(B)) = o(L/(UV)). \tag{65}$$

If it holds, Proposition 2.3 implies $g(p) \le 3q_p$ for every sufficiently large prime. The direct support threshold in Corollary 5.4 would also suffice. No converse is asserted. The weighted Hamming estimate and Proposition 3.3 provide only absolute entropy lower bounds of order $o(L)$; relative expansion and the folding estimates do not reach the target. In particular the very large common denominator needed for near-full rational entropy cannot be discarded by replacing the rational values with a single small-denominator model.

10.1. **An equivalent high-order common-divisor question.** To preserve the link to arithmetic matrices, we give the elementary equivalence directly. The detailed Smith-factor theory and the typical gcd distribution belong to the separate companion manuscript on arithmetic concentration, not to the main expansion theory here. For B as above put $P = \prod_{b\in B} b$. For $m = \lceil (UV)^2 \rceil$ take m independent copies of the difference variables C_b and their integer numerators $N_i = P\sum_b C_{bi}/b$. Let M be the $n \times (m-1)$ integer matrix with rows $(C_{b2} - C_{b1}, \dots, C_{bm} - C_{b1})$. Put $G = \gcd(N_2 - N_1, \dots, N_m - N_1)$ and count primes $p \in (X, 2X]$ dividing G only on the event that M has full rational column rank, setting the count Ω_X to zero otherwise. On full rank the vector of numerator differences is nonzero: a zero difference of integer numerators forces equality of all its digit coordinates after reduction at every b, which would contradict full rank. For such primes p this agrees with the prime count obtained using the scaled minors in the Smith formulation. Indeed $p \nmid P$ and a nonzero maximal minor has absolute value at most $(2\sqrt{m-1})^{m-1} = e^{o(L)} < p$. Hence the rows span modulo p. If $\mathbf{A}$ is the numerator-difference row, then all the determinants $\det(M_E, \mathbf{A})$ with $|E| = m - 2$ vanish

modulo p exactly when $\mathbf{A}$ is zero modulo p (extend a nonzero vector to a basis using rows of M). Dividing by $P_E = \prod_{b\in E} b$ preserves divisibility by p. These divisions produce integers by multilinearity and the formula $\mathbf{A} = \sum_b (P/b)\mathbf{v}_b$ with $\mathbf{v}_b$ a row of M.

Proposition 10.1 (Equivalence of the high-order gcd average)**.** *Put*

$$A_X = \max_{\substack{X<p\le 2X\\ p \text{ prime}}} \log(p\rho_p(B)).$$

For C sufficiently large,

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

In particular $\mathbb{E}\Omega_X \le X^{2-m}\exp(o(mL/(UV)))$ is equivalent to (65).

Proof. Let w_p denote the modular difference law; its nonnegative Fourier coefficients show $\|w_p\|_\infty = w_p(0) = \rho_p(B)$. Set $S_j(p) = \sum_x w_p(x)^j$. Equality of all m column numerators modulo p has probability $S_m(p)$. The rational-rank-deficient event implies rank deficiency modulo 2. A nonzero annihilator can be written as a vector $\lambda \in \mathbb{F}_2^m$ with coordinate sum zero, and thus has at least two nonzero coordinates. For each such vector fix two nonzero coordinates and condition on all the other columns. Equality modulo p of those other $m-2$ columns is necessary and has probability $S_{m-2}(p)$. At each row the two still-random entries have independent uniform parities, so annihilation of that row has probability $1/2$. A union bound over at most 2^m annihilators bounds the joint probability of deficiency and equality by $2^{m-n}S_{m-2}(p)$. Jensen applied with probability weights $w_p(x)$ gives $S_{m-2}(p) \le S_m(p)^{(m-3)/(m-1)}$, hence $S_{m-2}(p) \le p^2 S_m(p)$ because $S_m(p) \ge p^{1-m}$. As $n \sim (C/10)L$, increasing C makes the conditional rank-deficiency bound e^{-cL}. Thus

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

The point-mass bounds $\rho_p(B)^m \le S_m(p) \le \rho_p(B)^{m-1}$ give the result on taking logs: bound the sum by at most $2X$ times its largest upper bound and retain one prime attaining A_X for the lower bound. Since $L/m = o(L/(UV))$, the equivalence follows. □

The ordinary average $\sum_{X<p\le 2X} p\rho_p(B) \ll XU$ follows by counting prime divisors of a nonzero integer numerator of a signed reciprocal sum; its logarithmic size is $O(LU)$ and its zero probability is 2^{-n}. Thus Markov removes at most $O(XUe^{-aL/(UV)})$ primes with $\log(p\rho_p(B)) > aL/(UV)$ for any $a > 0$. If a decreases sufficiently slowly the omitted set is negligible in prime density. This ordinary average supplies no bound on the worst single prime, and increasing the moment order in itself does not prove (65).

10.2. **Three limitations.** First, a sparse rectangular block alone cannot provide full binary coverage. If D has $o(L)$ denominators, its sum has at most $2^{|D|} = p^{o(1)}$ possible values. More quantitatively, partition the torus of vectors $(ha^{-1}/p)_{a\in D}$ into $K^{|D|}$ cubes of side $1/K$. Pair counting within cubes gives

$$\#\{h \ne 0 : E_D(p,h) \le |D|/K^2\} \ge p/K^{|D|} - 1.$$

For a rectangular block of size $O(L/V)$, choosing $K = \lfloor V \rfloor$ gives $p^{1-o(1)}$ frequencies with energy $O(L/V^3) = o(L)$. Writing $N_D(t) = \#\{h \ne 0 : E_D(p,h) \le t\}$, the proposed full-range bound $\log(1+N_D(t)) \le 2t + o(L/(U^2V))$ for that sparse block would be false. Only the extreme low-energy region is controlled as in the spectral estimate of Section 2.1. A dense binary block is necessary for the entropy budget.

Second, short-relation exclusion and low moments alone are insufficient. For $n \asymp L$, $k_0 = O(L/U)$, the number of nonzero integer coefficient vectors c with $\sum_i |c_i| \le 2k_0$ is at most $K_0 = \binom{2n+2k_0}{2k_0} = e^{O(LV/U)}$. Fix the block H from Proposition 3.3, or any block of size $d = o(L)$ whose binary reciprocal subset sums are injective both rationally and modulo the chosen p. Rational dissociation alone would not suffice for the zero-coefficient diagonal. Choose w_i

independently in $[1, M_0]$ with $M_0 = \lceil 2 \cdot 3^d K_0 \rceil < p$. For any such nonzero coefficient vector and any fixed difference of H subset sums, the equality $\sum_i c_i w_i = \Delta$ modulo p occurs with probability at most $1/M_0$, by conditioning on all but one nonzero coefficient coordinate. A union bound gives a choice excluding all these relations. In particular its mixed moments through k_0 obey the same exact diagonal formula as those in (9), with the true prime weights replaced by w_i. Yet the integer sums on these artificial weights lie in $[0, nM_0]$, hence the sumset even after adding H and the sparse A has size at most $(nM_0 + 1)2^{|A|+d} = p^{o(1)}$. Rank statements obtained by conditioning on common values of column-wise maps do not distinguish these artificial weights from genuine reciprocal weights. Therefore abstract short-relation, rank, and low-moment information is not sufficient. This is not a counterexample with the true reciprocal coefficients.

Third, controlling joint divisibility for many large primes need not control isolated concentration. For a specified $p_0 \in (X, 2X]$ and fixed $a > 0$, take an artificial random integer $N = p_0 Z + W_0$, with independent Z uniform on $[0, \lfloor X^D \rfloor - 1]$ for fixed $D > 2$, and W_0 uniform on $[0, \lfloor p_0 e^{-aL/(UV)} \rfloor - 1]$. For every $q \le X^{D-1}$ coprime to p_0, its point probabilities modulo q equal $1/q + O(X^{-D})$ (count the progression of Z). Modulo p_0, however, it is uniform on only $\lfloor p_0 e^{-aL/(UV)} \rfloor$ points. Its mth peak moment at p_0 is $p_0^{1-m} \exp((a + o(1))mL/(UV))$, which alone violates the proposed high-order average. Any nonzero difference of its possible integer values has magnitude $O(X^{D+1})$, so after excluding the all-equal degeneracy the gcd can have only $O_D(1)$ primes in $(X, 2X]$. This shows the limitation of treating a factorial-moment bound for the number of prime factors as an estimate of the first moment. It is an abstract model, not a reciprocal-subset counterexample.

10.3. **Research direction.** The concrete remaining task is to exclude an isolated prime at which the product of cosine factors for the actual prime inverse vector is too concentrated. Section 3 supplies a quantitative arithmetic obstruction to exact low-rank generalized arithmetic progressions and a direct support-expansion theorem. One possible route is an exponential-scale inverse Littlewood–Offord argument to deduce a structural model within the rank–height budget from violation of (65). Section 6 avoids the general progression-cover step for a composite anchor at the larger height and obtains relative heterogeneous mixing, but neither it nor the preceding results supplies the required absolute near-full entropy estimate. A proof must have quantitative constants uniform at the scale $L/(UV)$; merely obtaining a small-rank progression with size $p^{1-o(1)}$ does not by itself exclude concentration. The theory of reciprocal sumsets and multilinear Kloosterman sums [6] is relevant, but no suitable uniform-in-moment-order estimate is invoked here. Results exhibiting nonuniformity of inverses on very short intervals [7] counsel against assuming independent random residues. Any future claim of all-primes coverage should establish a genuinely stronger pointwise bound (or the direct support criterion), not just an equivalent CRT, gcd, or matrix formulation.

Scope of the results

This paper concerns pointwise relative gains and rational information budgets. Uniform all-primes binary coverage at $O(\log p \log_2 p)$, and more strongly at $O(b(p))$, remains open. The sharp rational-budget theorem and its two threshold corollaries are new derivations from the recursive entropy input and the large-prime CRT mechanism. Their construction is based on distinct rational values, while the older Bernoulli constructions and window estimates are retained for their different probabilistic scope. No external mathematical refereeing, formal verification, or comprehensive priority search is claimed. Repeated sums and higher powers in proofs are auxiliary; where the conclusions concern unit fractions the denominator sets have been checked for distinctness.

References

[1] M. N. Bleicher and P. Erdős, *Denominators of Egyptian fractions II*, Illinois J. Math. **20** (1976), 598–613.

[2] P. Erdős and R. L. Graham, *Old and new problems and results in combinatorial number theory*, Monographies de L'Enseignement Mathématique 28, Geneva, 1980.
[3] H. Yokota, *On a problem of Bleicher and Erdős*, J. Number Theory **30** (1988), 198–207. doi:10.1016/0022-314X(88)90017-0.
[4] Y. P. Liu and M. Sawhney, *On further questions regarding unit fractions*, Int. Math. Res. Not. **2026**, no. 2, rnaf382. doi:10.1093/imrn/rnaf382. Numbering used here: arXiv:2404.07113v1, 10 April 2024.
[5] S. Bettin, L. Grenié, G. Molteni and C. Sanna, *A lower bound for the number of Egyptian fractions*, Math. Comp., published electronically 22 January 2026. doi:10.1090/mcom/4190.
[6] J. Bourgain and M. Z. Garaev, *Sumsets of reciprocals in prime fields and multilinear Kloosterman sums*, Izv. Math. **78** (2014), 656–707. doi:10.1070/IM2014v078n04ABEH002703.
[7] M. Z. Garaev and I. E. Shparlinski, *On the distribution of modular inverses from short intervals*, Mathematika **69** (2023), 1183–1194. doi:10.1112/mtk.12224.
[8] B. Green and I. Z. Ruzsa, *Freiman's theorem in an arbitrary abelian group*, J. London Math. Soc. (2) **75** (2007), 163–175. doi:10.1112/jlms/jdl021; arXiv:math/0505198.
[9] T. Tao and V. H. Vu, *Inverse Littlewood–Offord theorems and the condition number of random discrete matrices*, Ann. of Math. (2) **169** (2009), 595–632. `https://annals.math.princeton.edu/2009/169-2/p06`.
[10] H. H. Nguyen and V. H. Vu, *Optimal inverse Littlewood–Offord theorems*, Adv. Math. **226** (2011), 5298–5319. doi:10.1016/j.aim.2011.01.005.
[11] T. Sanders, *On the Bogolyubov–Ruzsa lemma*, Anal. PDE **5** (2012), 627–655. doi:10.2140/apde.2012.5.627; arXiv:1011.0107v2.
[12] G. Petridis, *Plünnecke's inequality*, Combin. Probab. Comput. **20** (2011), 921–938. Publisher record.
[13] C. Reiher and T. Schoen, *Note on the theorem of Balog, Szemerédi, and Gowers*, Combinatorica **44** (2024), 691–698. doi:10.1007/s00493-024-00092-5.
[14] D. Hajela and P. Seymour, *Counting points in hypercubes and convolution measure algebras*, Combinatorica **5** (1985), 205–214. doi:10.1007/BF02579363.
[15] H. J. Landau, B. F. Logan and L. A. Shepp, *An inequality conjectured by Hajela and Seymour arising in combinatorial geometry*, Combinatorica **5** (1985), 337–342.
[16] F. Gonçalves and D. Radchenko, *Sharp lower bounds for sumsets in hypercubes*, preprint, arXiv:2607.01458 (2026), Corollary 2. The classical Boolean specialization is the only inequality needed here.
[17] E. Kowalski and K. Soundararajan, *Equidistribution from the Chinese Remainder Theorem*, Adv. Math. **385** (2021), 107776. doi:10.1016/j.aim.2021.107776.
[18] P. Chen, *Recursive Entropy and Optimal Modular Coverage of Harmonic Subset Sums*, unpublished manuscript, 2026.