hexagon:2611.01202v1

Shifted-cofactor cancellation and a multiplicity threshold for product representations

Brice Pouly*

Abstract

Let $d \in (D, 2D]$ and $m \in (U, V]$ be independent with harmonic weights, and suppose the observer knows their product. Restrict d to have k prime factors counted with multiplicity, with $k/\log\log D \to r \in [0, 2)$. For any fixed integer $b \geq 1$, the sharp threshold for cancellation of $\lambda(m+b)$ on every window with $\log(V/U) \to \infty$ is the root $r_* = 0.542211\ldots$ of $(1+r)\log(1+r) - r\log r = 1$. Equality is included at every rate of approach. Cancellation holds uniformly in a test of the product, a divisor coefficient and a bounded multiplicative cofactor weight. It also holds for $\mu(m+b)$. If the observer additionally receives a message drawn from d alone, bounded conditional entropy of the message given the product still forces the optimal probability of predicting $\lambda(m+b)$ to tend to $1/2$. Below r_*, products with a unique eligible divisor have positive mass on sufficiently wide windows or when U exceeds a sufficiently large fixed power of $2D$. This holds on every widening window for $r < 0.481545\ldots$, and the recovery probability tends to one for $1 \leq k = o(\log\log D)$. The proof combines logarithmic two-point cancellation with a count of prime reallocations. At equality it uses the first favorable cutoff among finitely many scales.

1 Prediction from the product

Suppose one knows that $n = dm$, that $D < d \leq 2D$, and how many prime factors d has. When does n reveal the parity of the number of prime factors of $m+1$? If d is prime, a second prime in $(D, 2D]$ can divide n only by dividing m. Its probability is $O(1/\log D)$ under the sampling below. Usually one can therefore identify d from n, divide by it, and compute the sign. With more prime factors, some may move from d to m while a comparable product moves back. The exchange preserves the observation. The two shifted cofactors can have opposite signs.

For instance, with $D = 20$ and $9 < m \leq 15$, the product 360 admits $(d, m) = (24, 15)$ and $(36, 10)$. Both divisors have four prime factors counted with multiplicity. The two possible signs are $\lambda(16) = 1$ and $\lambda(11) = -1$. Under harmonic sampling the two representations have exactly the same weight. A second representation need not reverse the sign in general. The analytic step in the proof controls the combined effect of the different representations.

1.1 The experiment and the threshold

Write $\Omega(n) = \sum_p v_p(n)$, $\omega(n) = \sum_{p|n} 1$, $\lambda(n) = (-1)^{\Omega(n)}$, and let μ be the Möbius function. The parameters satisfy $D \to \infty$, $1 \leq U < V$, and $\ell = \log(V/U) \to \infty$. All logarithms are natural. For a nonempty set $\mathcal{D} \subseteq (D, 2D] \cap \mathbb{N}$ set

$$H(\mathcal{D}) = \sum_{d \in \mathcal{D}} \frac{1}{d}, \qquad W = \sum_{U < m \leq V} \frac{1}{m} = \ell + O(1).$$

*Research and writing with AI assistance.

The factors are independent and the pair probability is $1/(dmH(\mathcal{D})W)$. For an integer shift $b \geq 1$ and functions $f, g : \mathbb{N} \to \mathbb{C}$ bounded in modulus by one, define

$$\mathcal{N}_b(f, g; \mathcal{D}) = \sup_{\substack{|a_d| \leq 1 \\ |T(n)| \leq 1}} \left| \mathbb{E}\big[a_d g(m) f(m+b) T(dm)\big] \right|. \tag{1}$$

The coefficients and the product test may be complex and depend on all scales. Uniformity in g below refers to all functions bounded by one and multiplicative on coprime arguments; these functions may also depend on the scales. The exact normalization in (1) can be replaced by $H(\mathcal{D})\ell$ at relative cost $1 + O(1/\ell)$. For a layer

$$\mathcal{D}_k = \{D < d \leq 2D : \Omega(d) = k\}, \qquad H_k(D) = H(\mathcal{D}_k),$$

put

$$J(r) = (1+r)\log(1+r) - r\log r, \qquad J(0) = 0, \qquad J(r_*) = 1.$$

Since $J'(r) = \log(1 + 1/r) > 0$ for $r > 0$, the root is unique, and $r_* = 0.5422114197\ldots$. All limiting statements below hold along every sequence satisfying the displayed hypotheses; no relation between D and the cofactor window is imposed unless stated.

Theorem 1.1 (Cancellation threshold)**.** *Fix $b \geq 1$. If $k/\log\log D \to r \in [r_*, 2)$, then $\sup_g \mathcal{N}_b(f, g; \mathcal{D}_k) \to 0$ for $f = \lambda$ and $f = \mu$ on every widening window. This also holds when $\mathcal{D}$ consists of all integers in $(D, 2D]$. At $r = r_*$ the ratio $k/\log\log D$ may approach from either side at any rate.*

For the converse, we use a certificate available from n itself. Let

$$A_k(n) = \#\{d \in \mathcal{D}_k : d \mid n\}, \qquad M(n) = \#\{d \in \mathcal{D} : d \mid n,\ U < n/d \leq V\}.$$

On the layer, $A_k(n) = 1$ implies $M(n) = 1$ for every sampled n. The window affects the probability of obtaining this certificate.

Theorem 1.2 (Recovery below the threshold)**.** *Suppose $k/\log\log D \to r$ with $0 \leq r < r_*$ and $k \geq 1$. Set $B = 1 + \min(\log(2D), \log(2U))$.*

1. *If $\ell \geq B^\beta$ eventually for some fixed $J(r) < \beta < 1$, then $\liminf \mathbb{P}(A_k(dm) = 1) > 0$. The same holds, without the width condition, if $U \geq (2D)^{C_r}$ for a sufficiently large constant C_r depending only on r.*

2. *Define*

$$K(r) = \inf_{t>1}\{r\log(1+t) + 1/t\}, \qquad K(r_0) = 1, \qquad r_0 = 0.4815450284\ldots.$$

For $0 \leq r < r_0$, the same conclusion holds on every widening window. For $r = 0$, $\mathbb{P}(A_k(dm) = 1) \to 1$ on every such window.

3. *Under any of the preceding recovery conditions, for each fixed $b \geq 1$ and $f = \lambda, \mu$,*

$$\liminf \sup_{|T| \leq 1} |\mathbb{E}[f(m+b)T(dm)]| > 0.$$

At $r = 0$ this supremum tends to 1 for λ and $6/\pi^2$ for μ.

These results identify r_* as the exact threshold for cancellation to hold on *all* widening harmonic windows, within $0 \leq r < 2$. Below r_*, taking $U = 1$ makes B bounded, so part 1 supplies a family of windows where cancellation fails. Recovery on every individual window is proved here for $r < r_0$; for $r_0 \leq r < r_*$ it remains open outside the two regimes in part 1. The constant r_0 measures the reach of the second recovery estimate; its relation to an earlier optimization is explained in Section 6.

Layer limit	Conclusion	Cofactor windows
$r = 0$, $k \geq 1$	Unique cofactor with probability $1-o(1)$	Every $\ell \to \infty$
$0 < r < r_0$	Positive mass of unique recovery	Every $\ell \to \infty$
$r_0 \leq r < r_*$	Positive mass of unique recovery	The two regimes of Theorem 1.2(1)
$r_* \leq r < 2$	Uniform cancellation; $M(dm) \to \infty$ in probability	Every $\ell \to \infty$

Corollary 4.2 relates this multiplicity to the information needed to recover m from a list. The full dyad has the same cancellation and multiplicity conclusions. For an example in the recovery range, $U = D$ and $V = D\exp((\log D)^{\beta})$ in part 1 put both factors at size $D^{1+o(1)}$. All shifts are fixed, and the errors may depend on b; no global rate is asserted.

1.2 What the product reveals

The conditional law of the admissible representations of n is uniform, whereas

$$\mathbb{P}(dm = n) = \frac{M(n)}{nH(\mathcal{D})W}.$$

Thus the best probability of recovering m from n is $\mathbb{E}[1/M(dm)]$. For predicting the binary sign $\lambda(m+b)$, it is

$$\frac{1 + \mathbb{E}\,|\mathbb{E}(\lambda(m+b) \mid dm)|}{2}. \tag{2}$$

For fixed a_d, g, duality of the product test gives the exact identity

$$\begin{aligned} &\sup_{|T|\leq 1} |\mathbb{E}[a_d g(m) f(m+b) T(dm)]| \\ &\qquad = \frac{1}{H(\mathcal{D})W} \sum_n \frac{1}{n} \left| \sum_{\substack{d|n,\ d\in\mathcal{D} \\ U<n/d\leq V}} a_d g(n/d) f(n/d+b) \right|. \end{aligned} \tag{3}$$

With $a_d = g = 1$, this is the expectation of the absolute conditional mean in (2). For μ, which can vanish, Theorem 1.2(3) concerns correlation with a bounded test. Recovery and prediction place no restriction on the computational complexity of that test.

The shift matters to the information being lost. On an exact layer, the unshifted sign is always known:

$$\lambda(m) = (-1)^k \lambda(dm).$$

For the shifted value, the uniformity in a_d has an operational consequence. Suppose the observer also receives $d \bmod q$ for a fixed q. Throughout the cancellation range, the best guess of $\lambda(m+b)$ still succeeds with probability $1/2 + o(1)$. Proposition 4.3 allows any randomized message drawn from d alone whose conditional entropy given the product remains bounded.

1.3 Divisor problems and the proof mechanism

Classical multiplication-table problems take a random integer and ask whether it has a divisor in a prescribed interval. Tenenbaum studies the fraction H_1/H with exactly one such divisor among integers having at least one [20]. Ford determines the order of magnitude of H and, in a large range, of H_1 [4]. Maier and Tenenbaum prove the existence of close divisors for almost all integers [11]. More recently, Ford, Gabdullin and Pratt obtain estimates uniform in the number of required divisors in a dyadic interval [7]. Their counts use the natural measure on

integers. Here the observed product has mass $M(n)/(nH_kW)$, and its eligible divisors must also satisfy $\Omega(d) = k$.

Ford's extremal problem [5] asks how large one can make a set whose pairwise products are almost all distinct. The relation between few multiplicative collisions and random multiplicative sums is developed by Soundararajan and Xu [18]. In the present experiment the factor populations and their sampling weights are prescribed, and the shifted signs are deterministic arithmetic functions.

Schlitt's Theorem 1.2 and Corollary 1.3 [16] give a close comparison. Both factors in his multiplication table have all their primes in a common set of relative density δ satisfying his uniform prime-distribution hypothesis. The number of distinct products is comparable to the number of pairs for $\delta < 1/\log 4$, and is of smaller order at and above equality. The entropy of allocating primes between populations of sizes aL, bL is

$$\mathcal{E}(a, b) = (a + b)\log(a + b) - a\log a - b\log b.$$

The symmetric value $\mathcal{E}(\delta, \delta) = \delta \log 4$ and the asymmetric value $\mathcal{E}(r, 1) = J(r)$ explain the related constants. Here the factor-count restriction applies to the first factor alone, and the estimate is uniform against signed functions of the exact product.

Comparing prime allocations with the volume available to their logarithms is part of the multiplication-table method of Ford [4] and Koukoulopoulos [9, 10]; see in particular [10, Theorem 1.7 and Section 2.1]. Our allocations preserve an exact factor count on one factor. Their support enters an estimate for the signed conditional mean in (3).

In *Reconstruction and sign prediction from sampled prime factors* [15], the observations include sampled prime labels and, in Part III, a complete rough cofactor and the small-prime parts of three neighboring affine forms, distributed between two observers. The experiment here samples d and m independently under a prescribed law and reveals only their exact product. A related residue observation appears in *Parity information in a rough cofactor and exact reconstruction thresholds* [14]: the cofactor there is the quotient remaining after removal of all small prime powers. In *Observations and Shift Decoding in Divisor Incidences* [13], a list is used to locate a hidden displacement. The candidates in (3) are the equiprobable cofactors of a single product.

Here is how the comparison $J(r) = 1$ enters the proof. At a prime cutoff y, write $L = \log\log y$. If the two factors supply about rL and L prime occurrences, the number of ways of exchanging equal-sized sets of labels is bounded using

$$\sum_i \binom{a}{i}\binom{a'}{i} = \binom{a + a'}{a}, \qquad a \sim rL, \quad a' \sim L.$$

Its exponential size is $e^{J(r)L+o(L)}$. The range to be covered in the logarithm of a divisor has scale $\log y = e^L$. This comparison motivates the threshold. Different choices of labels can give the same divisor when a prime occurs repeatedly; the cancellation proof counts actual allocations.

For cancellation, we condition on the small-prime part of the product. Two allocations give distinct affine forms in the remaining cofactor, to which logarithmic two-point cancellation applies. The resulting square-root estimate is controlled by the support of the allocation count. Above r_*, typical profiles make that support small. At equality, a first favorable cutoff selected from the observed small primes supplies the missing gain. In the reverse direction, removing a largest exchanged prime gives a summable collision bound and an explicit window cost. A second count, using only the divisor profile, proves recovery on every widening window for $r < r_0$.

2 Arithmetic inputs and the local prime-factor law

2.1 A shift on the hidden factor

For a fixed bounded sequence $F : \mathbb{N} \to \mathbb{C}$, consider the condition

$$\frac{1}{\ell} \sum_{\substack{U<n\le V \\ n\equiv a \pmod q}} \frac{F(vn)\overline{F(v'n)}}{n} \longrightarrow 0 \qquad (v \neq v', \quad v, v', q, a \text{ fixed}) \tag{4}$$

on every widening window. All positive integers v, v', q and all residue classes a are allowed. For $F(n) = f(n+b)$ with fixed $b \geq 1$, the choice $f = \lambda$ satisfies (4) by Tao's logarithmic two-point theorem [19, Theorem 2]. For $f = \mu$, the multiplicative version applies [19, Corollary 5]: the prime values $\mu(p) = \lambda(p)$ give the same nonpretentiousness condition.

Indeed, writing $n = qh + a$ with $0 \leq a < q$ gives slopes qv, qv' and intercepts $av + b, av' + b$. Their determinant is $qb(v - v') \neq 0$. Replacing $1/(qh + a)$ by $1/(qh)$ for $h \geq 1$ costs $O_q(1)$ before division by ℓ. Rescaling the endpoints and discarding a bounded lower end changes the logarithmic length by $O_q(1)$. Tao's theorem permits that length to grow arbitrarily slowly.

The proof of Theorem 1.1 uses only (4) and $|F| \leq 1$, and yields the same conclusion for $g(m)F(m)$ under these assumptions. The sequence F is fixed throughout; the coefficients a_d, the multiplicative weight g, and the product test T need not be. Each arithmetic limit involves finitely many slopes. The nonzero shift is essential: with $b = 0$, $f = \lambda$, $a_d = \lambda(d)$, $g = 1$ and $T = \lambda$, the integrand in (1) equals one.

2.2 The small-prime parts on one layer

For $y \geq 3$, write

$$P_y = \prod_{p\le y} p, \qquad \mathcal{S}_y = \{n \in \mathbb{N} : p \mid n \Rightarrow p \leq y\},$$

including 1, and set

$$n_{\le y} = \prod_{p\le y} p^{v_p(n)}, \qquad \rho_z(y) = \prod_{p\le y}(1 - z/p) \quad (0 \leq z < 2).$$

At fixed y we abbreviate $\rho_z(y)$ to ρ_z. The Sathe–Selberg estimate for Ω gives, with $L_x = \log\log x$,

$$\#\{n \leq x : \Omega(n) = k\} = \frac{x}{\log x} \frac{L_x^{k-1}}{(k-1)!} \left\{ G\left(\frac{k-1}{L_x}\right) + O_\epsilon\left(\frac{k}{L_x^2}\right)\right\} \tag{5}$$

uniformly for $1 \leq k \leq (2 - \epsilon)L_x$, where

$$G(z) = \frac{1}{\Gamma(1+z)} \prod_p \frac{(1-1/p)^z}{1 - z/p}.$$

See Selberg [17, Theorem 4, pp. 85–87], and Tenenbaum [22, Chapter II.6, Section 6.1, pp. 299–304]. The product is positive and continuous on compact subintervals of $[0, 2)$. The local factor at 2 explains the restriction on the range.

Lemma 2.1 (Local law)**.** *Fix $y \geq 3$. Along every sequence with $D, \ell \to \infty$ and $k/\log\log D \to r \in (0, 2)$, the variables $u = d_{\le y}$ and $v = m_{\le y}$ converge jointly in total variation to independent variables on $\mathcal{S}_y$, with laws*

$$\mathbb{P}_y(u = u_0) = \rho_r \frac{r^{\Omega(u_0)}}{u_0}, \qquad \mathbb{P}_y(v = v_0) = \frac{\rho_1}{v_0}. \tag{6}$$

For every fixed $j \geq 0$ and $0 < A < B$,

$$\frac{1}{H_k(D)} \sum_{\substack{AD < d_0 \leq BD \\ \Omega(d_0)=k-j,\ (d_0,P_y)=1}} \frac{1}{d_0} \longrightarrow \rho_r r^j \frac{\log(B/A)}{\log 2}. \tag{7}$$

Proof. On compact ranges of x/D, replacing $\log x, L_x$ by $\log D, L_D$ in the density in (5) produces relative error $o(1)$, uniformly over such x. Indeed, $L_x - L_D = O(1/\log D)$ and $k = O(L_D)$. Partial summation over $(AD, BD]$ therefore multiplies the density by $\log(B/A)$. Lowering k by a fixed j multiplies the leading density by

$$\frac{(k-1)\cdots(k-j)}{L_D^j} \longrightarrow r^j.$$

Insert $\mathbf{1}_{(d_0,P_y)=1} = \sum_{h|(d_0,P_y)} \mu(h)$. For fixed squarefree h, the substitution $d_0 = he$ lowers the layer by a further $\Omega(h)$ and contributes $1/h$. The finite inclusion–exclusion sum gives ρ_r, proving (7).

Substituting $d = d_0 u_0$ gives the first law in (6). For $m = m_0 v_0$, periodicity of $(m_0, P_y) = 1$ gives restricted harmonic mass $\rho_1 \ell + O_{y,v_0}(1)$ before the factor $1/v_0$. The two Euler products in (6) each sum to one. Pointwise convergence on this countable space therefore gives total variation convergence by truncation. Independence follows from the sampling. A smaller cutoff is a deterministic projection of a larger one. □

3 Cancellation within product representations

Keep y fixed and write $d = d_0 u$, $m = m_0 v$, where $(d_0 m_0, P_y) = 1$. For fixed $d_0, m_0, w = uv$, each eligible $u \mid w$ yields the same observation $d_0 m_0 w$. The factors to be averaged, $F((w/u)m_0)$, have different slopes. This remains true after selecting allocations by their small-prime parts. That freedom will implement the stopping rule at equality.

The use of dilation correlations to control multiplicative weights is related to the criteria of Kátai [8] and Bourgain–Sarnak–Ziegler [2, Theorem 2]. Here the averaging takes place within factorizations of a fixed product, retaining a prescribed factor count and a dyadic constraint. The support estimate for these allocations determines the threshold. For $F(n) = \lambda(n+b)$ or $F(n) = \mu(n+b)$, with $T = 1$, uniformity in a bounded multiplicative cofactor weight is already supplied by Tao's nonasymptotic theorem [19, Theorem 3]; allowing arbitrary $T(dm)$ requires the product-fiber argument below.

For $E \subseteq \mathcal{S}_y^2$, let $\mathcal{N}(F; E)$ denote (1) with $f(m+b)$ replaced by $F(m)$, the indicator $\mathbf{1}_{(u,v)\in E}$ inserted, and an additional supremum over multiplicative g. For $w \in \mathcal{S}_y$, $j \geq 0$, $z > 0$, put

$$N_E(w, j, z) = \#\{u \mid w : (u, w/u) \in E,\ \Omega(u) = j,\ z < u \leq 2z\}, \qquad d\nu = \rho_r \rho_1 \frac{r^j}{w} \frac{dz}{z}.$$

The measure includes counting measure on w, j. Write $\mathbb{P}_y(E)$ for the product law in (6).

Proposition 3.1 (Finite-cutoff transfer)**.** *Fix $y \geq 3$, a set $E \subseteq \mathcal{S}_y^2$, and a complex sequence F with $|F| \leq 1$ satisfying (4). Along every sequence with $D, \ell \to \infty$ and $k/\log\log D \to r \in (0, 2)$,*

$$\limsup_{D,\ell\to\infty} \mathcal{N}(F; E) \leq \frac{1}{\log 2} \int \sqrt{N_E}\, d\nu, \tag{8}$$

$$\int N_E\, d\nu = (\log 2)\mathbb{P}_y(E). \tag{9}$$

In particular, the right side of (8) is at most $\sqrt{\mathbb{P}_y(E)\nu(N_E > 0)/\log 2}$. The parameters y, E, F are fixed before the arithmetic limit, while the supremum remains uniform in a_d, T, g.

Proof. First retain $u, v \le C$. By Lemma 2.1, the upper limit of the discarded unsigned probability tends to zero as $C \to \infty$. In a group indexed by $d_0, w, j = k - \Omega(d_0)$, give allocation u the coefficient $c_u = a_{d_0 u} g(w/u)$, of modulus at most one. The remaining multiplier $g(m_0)T(d_0 w m_0)$ is common to all allocations because $(m_0, v) = 1$. The reciprocal weight is $1/(d_0 w m_0)$. Before division by W, replacing $U/v < m_0 \le V/v$ by $U < m_0 \le V$ costs $O(1 + \log C)$ in each term, including when $U/v < 1$.

On the common interval put

$$\langle H \rangle = \frac{1}{W} \sum_{U < m_0 \le V} \frac{H(m_0)}{m_0}.$$

If A is the active set of allocations, Cauchy–Schwarz bounds the remaining average by

$$\langle \mathbf{1}_{(m_0, P_y)=1} \rangle^{1/2} \left(\sum_{u,u' \in A} c_u \overline{c_{u'}} \left\langle \mathbf{1}_{(m_0, P_y)=1} F((w/u)m_0) \overline{F((w/u')m_0)} \right\rangle \right)^{1/2}.$$

The first factor tends to $\sqrt{\rho_1}$. Diagonal entries in the second factor are at most $\rho_1 + o(1)$. Off-diagonal entries tend to zero by (4), summed over the finitely many reduced residue classes. The bound is therefore $\rho_1 \sqrt{\#A} + o_{y,C}(1)$, uniformly in d_0, in the active allocations and in their coefficients. Only finitely many slopes occur. Moreover, $D/C < d_0 \le 2D$ and $j \le \log C / \log 2$, so (7) bounds the total $1/d_0$ weight of the active layers by $O_{y,C}(H_k(D))$. This absorbs the endpoint and Gram-matrix errors even when $H_k(D)$ is small.

For fixed w, j, the truncated count $\#A$ is a step function of $z = D/d_0$. Apply (7) to its finitely many intervals of constancy; the $1/d_0$ weight gives Lebesgue measure in $\log z$. Summing over w, j and then letting $C \to \infty$ gives (8). Finally, a fixed allocation u is counted for $u/2 \le z < u$, an interval of logarithmic length $\log 2$. Opening the integral recovers the two point masses in (6), proving (9). The last bound follows by Cauchy–Schwarz in ν. □

The first moment of N_E has an exact integral. The remaining quantity is the measure of its support, without the multiplicity factor N_E. We estimate it first above the threshold and then at equality.

4 Above the threshold and at equality

Write $t = 1 + r$ and $L = \log \log y$. Under $\mathbb{P}_y$, the valuations of u have independent geometric laws of ratios r/p, and those of v have ratios $1/p$, independently. Consequently,

$$\mathbb{E}_y \Omega(u) = rL + O_r(1), \qquad \mathbb{E}_y \omega(uv) = tL + O_r(1),$$

and both variances are $O_r(1 + L)$. For the second estimate use $\mathbb{P}_y(p \mid uv) = t/p - r/p^2$. We shall also use

$$\sum_{w \in \mathcal{S}_y} \frac{t^{\omega(w)} \log(2w)}{w} \ll_r (\log y)^{t+1}. \tag{10}$$

Indeed, the product without the logarithm is $\prod_{p \le y}(1 + t/(p-1)) \ll_r (\log y)^t$. Under this normalized weight, $\mathbb{E} v_p(w) = O_r(1/p)$, so $\mathbb{E} \log(2w) = O_r(\log y)$. These estimates use Mertens' prime sum and partial summation. The same prime sum gives $\rho_1 \rho_r \asymp_r (\log y)^{-t}$.

4.1 Typical layers above equality

Fix $\eta > 0$ and retain pairs for which

$$|\Omega(u) - rL| \le \eta L, \qquad |\omega(uv) - tL| \le \eta L.$$

Their complement costs $o_{y\to\infty}(1)$ in the arithmetic upper limit, by Lemma 2.1 and Chebyshev's inequality. For fixed w, j, the logarithmic length of the set where $N_E(w,j,z) > 0$ is at most $\log(2w)$, since $1/2 \le z < w$. There are $O_r(1+L)$ possible j on the retained set, and on the support

$$r^j \le \exp\big((r\log r + O_r(\eta))L\big), \qquad 1 \le t^{\omega(w)-(t-\eta)L}.$$

Since $-t + r\log r - t\log t + (t+1) = 1 - J(r)$, (10) gives

$$\nu(N_E > 0) \ll_r (1+L)\exp\big((1 - J(r) + O_r(\eta))L\big). \tag{11}$$

When $J(r) > 1$, choose η sufficiently small and then send y to infinity in Proposition 3.1. This proves the strict upper range.

4.2 A first favorable cutoff at equality

At $r = r_*$, the main exponential factor in (11) is one. A single typical scale gives no decay. We use the freedom to choose the allocation scale from the small-prime data themselves. Realize all the laws $\mathbb{P}_y$ on one space of independent prime coordinates, and let u_y, v_y denote the products up to y. Set

$$B_y = \omega(u_y v_y)\log t - \Omega(u_y)\log r - \log\log y.$$

Its mean is $O_r(1)$ because $J(r) = 1$. Its variance is

$$\sigma^2 \log\log y + O_r(1), \qquad \sigma^2 = r(\log t - \log r)^2 + (\log t)^2 > 0. \tag{12}$$

A prime occurring once in u alone or in v alone gives increments $\log t - \log r$ and $\log t$, with probabilities $r/p + O_r(p^{-2})$ and $1/p + O_r(p^{-2})$. All other nonzero configurations contribute $O_r(p^{-2})$ to the second and third absolute moments together for large p, by geometric summation. The third absolute centered moment for one prime is therefore $O_r(1/p)$.

Lemma 4.1 (Favorable cutoffs)**.** *Put $L_s = 2^{3^s}$ and $y_s = \exp(\exp L_s)$ for $s \ge 1$. Almost surely, infinitely many s satisfy both*

$$B_{y_s} \ge 4\log L_s, \qquad \Omega(u_{y_s}) \le (r+1)L_s.$$

Proof. Chebyshev's inequality and the first Borel–Cantelli lemma give $B_{y_{s-1}}/\sqrt{L_s} \to 0$ almost surely, since $\sum_{s\ge 2} L_{s-1}/L_s < \infty$. The increments on successive disjoint prime blocks are independent. They have means $O_r(1)$, variances asymptotic to $\sigma^2 L_s$, and sums of third centered absolute moments $O_r(L_s)$. Lyapunov's central limit theorem makes the probability that a block increment exceeds $\sigma\sqrt{L_s}/2$ bounded below for large s. The second Borel–Cantelli lemma gives infinitely many such blocks. Their increments dominate $4\log L_s$ and the preceding contribution. The count condition fails only finitely often by Chebyshev's inequality, because $\sum_s 1/L_s < \infty$. □

Fix integers $s_0 \le S$ first. Partition the pairs by the *first* index $s \in [s_0, S]$ satisfying the two conditions of Lemma 4.1, together with the residual event of no such index. Membership in the stratum indexed by s depends only on (u_{y_s}, v_{y_s}), since this pair determines all earlier cutoffs. Proposition 3.1 therefore applies at y_s. A selection that used later prime coordinates would not have this property.

On the support at scale s,

$$r^j \le t^{\omega(w)} e^{-L_s} L_s^{-4}.$$

There are $O_r(L_s)$ choices of j. The support-length bound and (10) give

$$\nu(N_E > 0) \ll_r e^{-tL_s} L_s e^{-L_s} L_s^{-4} e^{(t+1)L_s} \ll_r L_s^{-3}.$$

Thus the restricted norm is $O_r(L_s^{-3/2})$. The remaining pairs cost at most their unsigned mass. At fixed s_0, S, Lemma 2.1 identifies its limit with the no-hit probability in the independent model. The triangle inequality gives

$$\limsup_{D,\ell\to\infty} \sup_{a_d,T,g} |\mathbb{E}[a_d g(m)F(m)T(dm)]| \le C_r \sum_{s\ge s_0} L_s^{-3/2} + \mathbb{P}(\text{no hit from } s_0 \text{ through } S). \quad (13)$$

Given $\epsilon > 0$, first choose s_0 so that C_r times the series tail is less than $\epsilon/3$. Lemma 4.1 then supplies a finite $S \ge s_0$ for which the no-hit probability is less than $\epsilon/3$. Keep these cutoffs fixed before taking the arithmetic limit. Lemma 2.1 and Proposition 3.1 make the full norm less than ϵ for all sufficiently large parameters along the chosen sequence. Only convergence of $k/\log\log D$ to r_* has been used.

The positive recovery bounds for each fixed $r < r_*$ are not uniform as $r \uparrow r_*$. They are therefore consistent with cancellation when the layer ratios converge to r_* from below.

4.3 Removing the layer and counting the representations

Under harmonic sampling of all $D < d \le 2D$, one has $\Omega(d)/\log\log D \to 1$ in probability. To see this, first use uniform counting on the dyad. The sum of $\mathbf{1}_{p|d}$ over $p \le D^{1/4}$ has mean $\log\log D + O(1)$ and variance $O(\log\log D)$: one- and two-prime counts have error $O(1/D)$ per choice. Repeated powers have bounded expected total contribution because $\sum_p \sum_{a\ge 2} p^{-a} < \infty$. The number of prime factors exceeding $D^{1/4}$ is bounded. Harmonic weights on the dyad are comparable to uniform counting. Consequently all but $o(1)$ of their mass lies in

$$|k - \log\log D| \le (\log\log D)^{2/3}.$$

The norm of a mixture is at most the mixture of the norms. The bound for $r = 1$ tends to zero uniformly in these layers, since a violating sequence would itself satisfy $k/\log\log D \to 1$. This proves the full-dyad assertion of Theorem 1.1.

Additional information and recovery. Write $N = dm$. We first ask how much information would allow a list of s guesses to contain m. We then ask whether a message drawn from the divisor can help predict a single sign. Entropies and mutual information below use natural logarithms.

Corollary 4.2 (Lists of candidate cofactors)**.** *In the full dyad and the layers of Theorem 1.1, $M(N) \to \infty$ in probability. For an integer $s \ge 1$, let $\delta_s = \mathbb{E}\min(1, s/M(N))$. Even with an additional discrete observation Z drawn from the pair (d, m), the success probability p of any list of at most s candidates based on (N, Z) satisfies*

$$p\log(1/\delta_s) \le \mathsf{I}(m; Z \mid N) + \log 2.$$

In particular $p \to 0$ for each fixed s if the conditional mutual information stays bounded. For a list guaranteed to contain m, the least possible conditional mutual information is exactly $\mathbb{E}\log^+(M(N)/s)$, where $\log^+ x = \max(0, \log x)$.

Proof. Take independent ξ_j uniform on $\{1, -1, i, -i\}$. For $S = \sum_j \xi_j u_j$ and $A = \sum_j |u_j|^2$, expansion gives $\mathbb{E}_\xi|S|^2 = A$ and $\mathbb{E}_\xi|S|^4 = 2A^2 - \sum_j |u_j|^4$. Hölder's inequality yields $\mathbb{E}_\xi|S| \ge \sqrt{A/2}$. Apply this with $g = 1$, $f = \lambda$ and $a_d = \xi_d$ in (3). On a fiber of size M, the M normalized summands have modulus $1/M$. Averaging over the phases gives $\mathbb{E}[M(N)^{-1/2}] \le \sqrt{2}\,\mathcal{N}_b(\lambda, 1; \mathcal{D})$, which tends to zero. The same phase may occur on many fibers because only expectation is being taken.

Without Z, δ_s is the optimal list success by uniformity on each fiber. For a list based on (N, Z), compare the true law P with the law Q that preserves (N, m) and (N, Z) separately

and makes m and Z conditionally independent given N. Under Q the success event has probability $q \le \delta_s$. The log-sum inequality applied to success and failure [3, Chapter 2] gives

$$\mathsf{I}(m; Z \mid N) = D(P\|Q) \ge p \log \frac{p}{q} + (1-p) \log \frac{1-p}{1-q} \ge p \log(1/\delta_s) - \log 2$$

with the usual limiting conventions. For fixed s, $\delta_s \to 0$ follows from divergence of $M(N)$. If success is certain, then $\mathsf{H}(m \mid N, Z) \le \mathbb{E} \log \min(s, M(N))$, whereas $\mathsf{H}(m \mid N) = \mathbb{E} \log M(N)$. Equality is attained by reporting a uniformly chosen subset of size $\min(s, M(N))$ of the fiber that contains the sampled cofactor. Conditional on that subset and N, its members are equiprobable. □

Proposition 4.3 (A message from the divisor)**.** *Suppose the conditional law of a countable-valued message Z given (d, m) depends only on d. Let $p_{(j)}(n)$ be the probabilities of the messages given $N = n$, arranged in decreasing order and padded with zeros if necessary. For every integer $L \ge 1$, all coefficients and weights admitted in* (1)*, and $|T| \le 1$,*

$$|\mathbb{E}[a_d g(m) f(m+b) T(N, Z)]| \le \sqrt{2L}\, \mathcal{N}_b(f, g; \mathcal{D}) + \mathbb{E}\left[\sum_{j>L} p_{(j)}(N)\right].$$

In the regimes of Theorem 1.1, $\mathsf{H}(Z \mid N) = O(1)$ therefore suffices for cancellation with this additional observation, including at $r = r_*$.

Proof. Write $K_d(j) = \mathbb{P}(Z = j \mid d)$, $Y = a_d g(m) f(m+b)$, and $c_j(n) = \mathbb{E}[Y K_d(j) \mid N = n]$. Product-test duality with the additional message gives the upper bound $\mathbb{E} \sum_j |c_j(N)|$. Take independent phases ξ_j as above. The coefficient $a_d \sum_j \xi_j K_d(j)$ has modulus at most one and depends only on d. Optimizing the product test for each realization and then averaging over the phases gives

$$\mathbb{E}\Big(\sum_j |c_j(N)|^2\Big)^{1/2} \le \sqrt{2}\, \mathbb{E}_\xi \mathbb{E}\Big|\sum_j \xi_j c_j(N)\Big| \le \sqrt{2}\, \mathcal{N}_b(f, g; \mathcal{D}).$$

Countable sums follow by truncation since $\sum_j |c_j(n)| \le 1$. On each fiber, sum over its L most likely messages by Cauchy–Schwarz. The remaining terms satisfy $|c_j(n)| \le \mathbb{P}(Z = j \mid N = n)$. This proves the bound even when the likely messages vary with n. Every probability after rank L is at most $1/(L+1)$, so their expected mass is at most $\mathsf{H}(Z \mid N)/\log(L+1)$. First keep L fixed in the arithmetic limit and then let L increase. □

For λ the optimal success probability based on (N, Z) is consequently $1/2 + o(1)$ under the proposition's hypotheses. A joint observer can instead send the one bit $Z = \lambda(m+b)$. This reveals the sign exactly, while Corollary 4.2 still rules out recovery of m with any fixed list, since $\mathsf{I}(m; Z \mid N) \le \log 2$.

A quartic test. With $C(n) = \lambda(n)\lambda(n+1)\lambda(3n+1)\lambda(3n+2)$, take $T = C$, $a_d = 1$ and $g = \lambda$. In the cancellation regimes,

$$\mathbb{E}\big[C(dm)\mathbf{1}_{\lambda(m)=\lambda(m+1)}\big] = \tfrac{1}{2}\mathbb{E}[C(dm)] + o(1).$$

The disagreement class satisfies the same formula. The two classes share the signed mass equally, for any value of $\mathbb{E}[C(dm)]$. The fixed-form correlation in *An explicit bound below one half for a quartic Liouville correlation* [12] is a harmonic average over a single origin. Here that origin is weighted by its number of admissible representations.

5 What a competing divisor would require

Two divisors in the same dyad are comparable. If $d' \neq d$, $d' \in \mathcal{D}_k$ and $d' \mid dm$, divide d, d' by their greatest common divisor. The resulting integers x, y satisfy

$$x \mid d, \quad y \mid m, \quad x/2 < y < 2x, \quad (x,y) = 1, \quad \Omega(x) = \Omega(y) \geq 1. \tag{14}$$

Indeed $y \mid xm$ and $(x,y) = 1$, and the common part cancels out of the two prime-factor counts. We will rule out even this necessary condition on a positive proportion of pairs. The collision estimate is independent of the target sign.

For integers $j \geq 1$, put $Z_j = 2\exp(\exp(j+1))$ and $\Omega_z(n) = \sum_{p \leq z} v_p(n)$. Fix an integer $h \geq 1$, let $Y_h = \exp(\exp h)$, and use the event

$$R_h = \{(m, Q_h) = 1\}, \qquad Q_h = \prod_{p \leq 2Y_h} p, \qquad \rho(h) = \varphi(Q_h)/Q_h.$$

Its probability tends to $\rho(h) > 0$ for fixed h. On R_h, a solution to (14) has $y > 2Y_h$ and $x > Y_h$. Hence $j = \lfloor \log\log x \rfloor \geq h$, and every prime of x, y is at most Z_j.

Lemma 5.1 (Prime counts at all exchange scales)**.** *Let $0 \leq r < r_*$, $k \geq 1$, and $k/\log\log D \to r$. For every fixed $\eta > 0$, there are $c, C > 0$, depending only on r, η, such that for every sufficiently large fixed h,*

$$\limsup_{D,\ell \to \infty} \mathbb{P}\{\exists j \geq h : \Omega_{Z_j}(d) > (r+\eta)j\} \leq Ce^{-ch}, \tag{15}$$

$$\limsup_{D,\ell \to \infty} \mathbb{P}\{\exists j \geq h : \Omega_{Z_j}(m) > (1+\eta)j \mid R_h\} \leq Ce^{-ch}. \tag{16}$$

The same bound as in (16) *holds without conditioning.*

Proof. Put $L_D = \log\log D$. Choose $0 < \gamma < 1$ with $r < (r+\eta)\gamma$. For $j \geq \gamma L_D$, the first upper cap follows from $\Omega(d) = k$ for all sufficiently large D. For $j < \gamma L_D$, if s has only primes up to Z_j and $a = \Omega(s) \leq k$, then $s \leq Z_j^k = D^{o(1)}$, uniformly. For $1 \leq a < k$, (5) on dyads gives

$$\mathbb{P}(s \mid d) = \frac{H_{k-a}(D/s)}{sH_k(D)} \ll_{r,\eta} \frac{(k/L_D)^a}{s}. \tag{17}$$

The function G is bounded above and away from zero over the relevant range, $\log\log(D/s) \leq L_D$, and the factorial ratio is at most k^a. Since $D/s = D^{1-o(1)}$ uniformly, partial summation on the two dyads gives uniform upper and lower bounds by constant multiples of their main terms. For $a = k$, the only possible value would be $d = s < D$; for $a > k$, there is no multiple on the layer. This also covers bounded k.

For $a_0 > 1$ let b_{a_0} be multiplicative, with $b_{a_0}(1) = 1$ and $b_{a_0}(p^i) = (a_0 - 1)a_0^{i-1}$ for $i \geq 1$. The positive expansion

$$a_0^{\Omega_z(n)} = \sum_{\substack{s \mid n \\ p \mid s \Rightarrow p \leq z}} b_{a_0}(s)$$

allows (17) to be used term by term. Take a small $\epsilon > 0$ and then $a_0 > 1$ close to one, so that $a_0(r+\epsilon) < 2$. With $z = Z_j$, the expected value for d is

$$\ll \prod_{p \leq Z_j} \left(1 + \frac{(a_0-1)(r+\epsilon)}{p - a_0(r+\epsilon)}\right) \ll \exp((a_0-1)(r+\epsilon)j).$$

Here $\log\log Z_j = j + O(1)$. Choose $\epsilon < \eta$ and a_0 sufficiently close to one that $(a_0-1)(r+\epsilon) < (r+\eta)\log a_0$. Markov's inequality and a geometric series over j prove (15).

For fixed Q, periodicity gives

$$\sum_{\substack{A<n\le B'\\(n,Q)=1}}\frac{1}{n}=\frac{\varphi(Q)}{Q}\log^+\frac{B'}{\max(1,A)}+O_Q(1)\qquad(0<A<B'),\tag{18}$$

where $\log^+ x=\max(0,\log x)$. Therefore, uniformly for positive integers s,

$$\mathbb{P}(\{s\mid m\}\mid R_h)\le\frac{1+O_h(1/\ell)}{s}\quad\text{if }(s,Q_h)=1,\tag{19}$$

and the probability is zero otherwise. Applying the same positive expansion with ratio 1 instead of $r+\epsilon$ gives a bound $\ll(1+O_h(1/\ell))\exp((a_0-1)j)$. The factor depending on h,ℓ multiplies the entire Euler product once. Taking $1<a_0<2$ near one proves (16). Setting $Q=1$ gives the unconditioned version. □

5.1 Counting both sides of an exchange

We first prove Theorem 1.2(1) for $r>0$. Fix $J(r)<\alpha<\beta<1$. Condition on R_h, discard the two exceptional sets in Lemma 5.1, and fix an admissible d. The variables d and R_h are independent. In an exchange at scale j, expose the largest prime $q\mid y$ and put $m'=m/q$, $y'=y/q$. Since $y>x/2$ and $\Omega(y)\le(1+\eta)j$,

$$q\le Z_j,\qquad\log q\gg_\eta e^j/j,\qquad\frac{x}{2y'}<q<\frac{2x}{y'}.\tag{20}$$

The last interval has endpoint ratio 4. Chebyshev's upper bound for the number of primes shows that the sum of $1/q$ over its primes satisfying (20) is $O_\eta(je^{-j})$.

Here is the union bound before changing the cofactor window. Let

$$\mathcal{G}_h=\{m:\Omega_{Z_j}(m)\le(1+\eta)j\text{ for every }j\ge h\},\qquad W_h=\sum_{\substack{U<m\le V\\(m,Q_h)=1}}\frac{1}{m}.$$

Choose $c_\eta>0$ such that every retained exchange satisfies $\log q\ge c_\eta e^j/j$. For the fixed admissible d, define

$$\Sigma_j(d)=\sum_{\substack{m'\ge1,(m',Q_h)=1\\\Omega_{Z_j}(m')\le(1+\eta)j}}\frac{1}{m'}\sum_{\substack{x\mid d\\e^j\le\log x<e^{j+1}}}\sum_{\substack{y'\mid m',\ y'\in\mathcal{S}_{Z_j}\\\Omega(y')=\Omega(x)-1}}\sum_{\substack{q\text{ prime},\ q\le Z_j\\\log q\ge c_\eta e^j/j\\x/(2y')<q<2x/y'\\U<qm'\le V}}\frac{1}{q}.\tag{21}$$

Every retained collision supplies at least one term of this sum, with the exact weight $1/(qm')$. Dropping coprimality and the requirement that q be the largest prime only adds terms. Hence

$$\mathbb{P}(A_k(dm)>1,\ m\in\mathcal{G}_h\mid d,R_h)\le\frac{1}{W_h}\sum_{j\ge h}\Sigma_j(d).$$

For each m' in (21), put $a=\Omega_{Z_j}(d)\le(r+\eta)j$ and $a'=\Omega_{Z_j}(m')\le(1+\eta)j$. Before choosing q, the number of pairs of divisors x,y' is at most

$$\sum_{i\ge1}\binom{a}{i}\binom{a'}{i-1}=\binom{a+a'}{a'+1}.\tag{22}$$

To obtain this bound, label the prime occurrences, including repetitions. Every pair of actual divisors is represented by at least one pair of label sets, and $i = \Omega(x) = 1 + \Omega(y')$. The sum is nondecreasing in both capacities. Stirling's formula at their upper endpoints gives

$$\binom{a+a'}{a'+1} \le (1+j)^{O_r(1)} \exp\bigl(\mathcal{E}(r+\eta, 1+\eta)j\bigr).$$

Choose η sufficiently small. Continuity at $(r,1)$ and absorption of the polynomial and the factor j in (20) bound the aggregate reciprocal q-weight by $C_{r,\alpha}e^{-(1-\alpha)j}$ for every such m'.

The condition $U < qm' \le V$, with $q \le Z_j$, implies $U/Z_j < m' \le V$. The bound just obtained is uniform on all admitted m', so nonnegativity permits us to enlarge their range and remove the profile condition:

$$\Sigma_j(d) \le C_{r,\alpha}e^{-(1-\alpha)j} \sum_{\substack{U/Z_j < m' \le V \\ (m', Q_h)=1}} \frac{1}{m'}.$$

By (18), this final reciprocal mass is at most

$$\rho(h)\bigl(\ell + \min(\log U, \log Z_j)\bigr) + O_h(1), \qquad W_h = \rho(h)\ell + O_h(1).$$

The lower endpoint is responsible for the extra term. Since $x \le 2D$, relevant shells have $e^j \le \log(2D)$, while $\log Z_j \ll e^j$. Thus the extra length is $O(\min(e^j, B))$, with B as in Theorem 1.2. Summing the shells bounds the conditional collision probability on retained pairs by

$$O_{r,\alpha}\left(e^{-(1-\alpha)h} + \frac{B^\alpha}{\ell}\right) + o_h(1). \tag{23}$$

Indeed, splitting the sum at $\log B$ gives

$$\sum_{j \ge 0} e^{-(1-\alpha)j} \min(e^j, B) \ll_\alpha B^\alpha.$$

The term B^α/ℓ tends to zero under the width hypothesis, including when B stays bounded. Together with Lemma 5.1, this proves

$$\lim_{h\to\infty} \limsup_{D,\ell\to\infty} \mathbb{P}(A_k(dm) > 1 \mid R_h) = 0. \tag{24}$$

One first fixes a large h and then uses the positive limiting mass $\rho(h)$. The strength of this conditional estimate will also give the Möbius assertion in Section 7.

5.2 Windows above a sufficiently large power of the divisor scale

When m is much larger than d, all competitors can be tested using comparatively small primes of m. This gives a second way around the boundary cost in (23).

We use the classical Kubilius approximation in total variation; see Tenenbaum [21, pp. 1099–1101]. A direct vector formulation is Ford [6, Lemma 2.1]: for m uniform on the integers up to X, the vector $(v_p(m))_{p\le z}$ and independent geometric variables with probabilities $(1-1/p)p^{-a}$ satisfy

$$\mathrm{d_{TV}} \ll_\epsilon u^{-u} + X^{-1+\epsilon}, \qquad u = \frac{\log X}{\log z}, \qquad 2 \le z \le X.$$

Taking $\epsilon = 1/2$ implies the weaker bound Ce^{-cu} with absolute positive constants. This is the only strength needed here. In Tenenbaum's additive formulation, $f(m) = \log m$ encodes the same vector through $\log(m_{\le z})$.

Let G be any indicator depending on this vector and let $\pi_z(G)$ be its independent-model mean. For $X \ge U \ge z$, the prefix sum of $G - \pi_z(G)$ is $O(X\varepsilon)$, where $\varepsilon = Ce^{-c\log U/\log z}$, allowing the integer endpoint. Partial summation gives the displayed harmonic transfer

$$\left| \frac{1}{W} \sum_{U < m \le V} \frac{G(m)}{m} - \pi_z(G) \right| \ll e^{-c\log U/\log z} \left(1 + \frac{1}{\ell}\right), \qquad 2 \le z \le U. \tag{25}$$

Its constant is uniform in G and in the window. It remains valid after averaging an event depending on a separately sampled d.

Take $z = 2D$. For fixed d, the event $A_k(dm) > 1$ is measurable in the vector, because each rival $d' \le 2D$ is tested by $d'/(d,d') \mid m$. For fixed h, R_h is also measurable there when D is sufficiently large. To determine their independent-model intersection probability, compare with an auxiliary window starting at

$$U' = \exp((\log(2D))^2), \qquad \log(V'/U') = (\log(2D))^2.$$

The approximation error tends to zero there, and (24) applies. Fix h large enough that the upper limiting collision probability conditional on R_h in this auxiliary window is less than $c_{\mathrm{sf}}/8$, where the absolute constant $c_{\mathrm{sf}} > 0$ is defined in (30). The upper limit of the joint probability in the independent model is then at most $\rho(h)c_{\mathrm{sf}}/8$.

With this h fixed, choose C_r large enough that (25) on every target window with $U \ge (2D)^{C_r}$ is eventually less than $\rho(h)c_{\mathrm{sf}}/8$. Since $\mathbb{P}(R_h) \to \rho(h)$ by periodicity,

$$\limsup_{D,\ell\to\infty} \mathbb{P}(A_k(dm) > 1 \mid R_h) < c_{\mathrm{sf}}/2 \qquad \big(U \ge (2D)^{C_r},\ \ell \to \infty\big). \tag{26}$$

The approximation error need only be this small fixed constant; it need not tend to zero at fixed C_r. This proves positive mass of unique recovery and suffices for the squarefree restriction in Section 7. For $0 < r < r_*$, the choices depend only on r, since the intermediate exponents in (23) may be fixed between $J(r)$ and 1. The case $r = 0$ is treated directly below.

6 Recovery on every widening window

The cost B^α/ℓ arose from summing over $m' = m/q$ after deleting a large exchanged prime. For smaller layers, we can afford to sum over all candidate y before looking at m. The advantage is that $\mathbb{P}(\{y \mid m\} \mid R_h)$ has the uniform majorant (19) on every widening window. The price is that the count now uses a typical prime profile only for d.

Lemma 6.1 (An exponentially weighted count)**.** *For fixed $0 < z < 1$ and $x \ge 3$,*

$$\sum_{x/2 < y < 2x} \frac{z^{\Omega(y)}}{y} \ll_z (\log x)^{z-1}.$$

Proof. It suffices to bound the ordinary sum up to $X = 2x$ and divide by $x/2$. For $1 \le i \le L_X = \log\log X$, summing the upper bound in (5) gives

$$\sum_{i \le L_X} z^i \#\{n \le X : \Omega(n) = i\} \ll \frac{X}{\log X}\, z \sum_{i \ge 1} \frac{(zL_X)^{i-1}}{(i-1)!} \ll_z X(\log X)^{z-1}.$$

For $\Omega(n) > L_X$, one has $z^{\Omega(n)} \le e^{L_X \log z} \le (\log X)^{z-1}$. The term $n = 1$ and bounded X are harmless. This estimate is also a special case of the multiplicative mean-value bound used by Maier and Tenenbaum [11, Lemma 1]. □

Fix $t > 1$. For a divisor $x \mid d$ in shell $j = \lfloor \log\log x \rfloor$, equality of factor counts in (14) implies

$$\sum_{\substack{x/2<y<2x \\ \Omega(y)=\Omega(x)}} \frac{1}{y} \le t^{\Omega(x)} \sum_{x/2<y<2x} \frac{t^{-\Omega(y)}}{y} \ll_t t^{\Omega(x)} e^{-(1-1/t)j}. \tag{27}$$

To sum over x, use

$$\sum_{\substack{x \mid d \\ p \mid x \Rightarrow p \le Z_j}} t^{\Omega(x)} = \prod_{\substack{p^a \| d \\ p \le Z_j}} (1 + t + \cdots + t^a) \le (1+t)^{\Omega_{Z_j}(d)}.$$

For any d satisfying the upper cap of Lemma 5.1 and every $j \ge h$, the sum over all possible exchanges, with weight $1/y$, is at most

$$C_t \exp\big(((r+\eta)\log(1+t) + 1/t - 1)j\big). \tag{28}$$

This is a union bound. It allows repetitions of the same y and has already discarded coprimality and all constraints on m/y.

The derivative in t of $r\log(1+t) + 1/t$ vanishes at $rt^2 = 1 + t$. For $0 < r < 2$, the unique solution in $t > 1$ is

$$t_r = \frac{1 + \sqrt{1+4r}}{2r},$$

and it gives the minimum defining $K(r)$. The function K is continuous and strictly increasing there, because its derivative is $\log(1 + t_r) > 0$. Also $K(0) = 0$ by letting $t \to \infty$. Substitution gives $K(0.48) < 1 < K(0.49)$, while $J(0.5) < 1$, so $r_0 < r_*$.

The same auxiliary constant occurs in Bado's optimization of a 2-adic bound for the divisibility problem $\varphi(n) \mid n + a$ [1, Section 7, Proposition 5]. In his notation,

$$H_2(\rho) = \frac{\rho}{2-\rho}, \qquad I_2(r) = \sup_{0<\rho<1} \{1 - H_2(\rho) + r\log\rho\}.$$

The substitution $\rho = 2/(1+t)$ is a bijection from $t > 1$ to $0 < \rho < 1$ and gives

$$I_2(r) - r\log 2 = \sup_{t>1}\{1 - 1/t - r\log(1+t)\} = 1 - K(r). \tag{29}$$

Thus his crossing equation $I_2(r) = r\log 2$ yields exactly r_0. Here the optimization controls recovery in the prescribed product-sampling experiment; the proof uses the estimate (28) derived above.

If $r < r_0$, first choose $t > 1$ with $r\log(1+t) + 1/t < 1$, and then choose $\eta > 0$ so the exponent in (28) remains negative. Conditional on R_h, every exchange has $j \ge h$. By (19), the probability of any such exchange is bounded by the sum of (28) over these shells, up to the factor $1 + O_h(1/\ell)$. Exceptional d cost $O(e^{-ch})$ by (15). Hence (24) holds on every widening window in this range. This proves the positive recovery assertion in Theorem 1.2(2).

Sparse layers. At $r = 0$, the positive-mass restriction on m can be removed. For each fixed h, the probability that d has a prime divisor at most $2Y_h$ tends to zero. For $k = 1$, this is immediate. For $k \ge 2$, apply (17) with $s = p$, or its proof at a fixed p, and sum to obtain $O_h(k/L_D) = o(1)$. On the remaining d, every exchange has $x > Y_h$, since $x \mid d$ and $x > 1$. Use (27)–(28) with the unconditional bound

$$\mathbb{P}(y \mid m) \le \frac{1 + O(1/\ell)}{y},$$

which follows from (18) with $Q = 1$. After (15), the upper limit of the collision probability is $O(e^{-ch})$. Sending $h \to \infty$ proves $\mathbb{P}(A_k(dm) = 1) \to 1$.

For the prime layer, the introductory estimate follows directly from $\sum_{D<p\le 2D} 1/p \ll 1/\log D$ and the same divisibility bound. The conditional bound is independent of the selected prime, so the conclusion also holds if its selection weight $1/p$ is replaced by $(\log p)/p$. The sparse-layer argument permits growing layers such as $k = \lfloor (\log\log D)/\log\log\log D \rfloor$.

7 The sign recovered on a unique representation

If $A_k(n) = 1$, its unique divisor and the associated cofactor can be computed from n. Define $T(n) = f(n/d + b)$ on this event and $T(n) = 0$ elsewhere. For $f = \lambda$, the correlation is exactly $\mathbb{P}(A_k(dm) = 1)$. For $f = \mu$, it is the mass of the same event with $m + b$ squarefree. Positive masses for the two events separately would not suffice. The conditional collision estimate (24) puts them together.

Lemma 7.1. *For fixed squarefree Q and $b \ge 1$, the conditional probability that $m + b$ is squarefree given $(m, Q) = 1$ tends, on every widening harmonic window, to*

$$\prod_{p\nmid Q}\Big(1-\frac{1}{p^2}\Big)\prod_{\substack{p\mid Q\\ p\nmid b}}\Big(1-\frac{1}{p(p-1)}\Big) \ge c_{\mathrm{sf}} := \prod_{p}\Big(1-\frac{1}{p(p-1)}\Big) > 0. \tag{30}$$

The lower bound is independent of Q, b.

Proof. For natural density, impose the square exclusions for finitely many primes and apply the Chinese remainder theorem together with $(m, Q) = 1$. If $p \mid Q$ and $p \nmid b$, precisely one of the $p(p-1)$ units modulo p^2 is forbidden. If $p \mid (Q, b)$, none is forbidden. If $p \nmid Q$, one of the p^2 classes is forbidden. For $m \le X$, the proportion excluded by a further square p^2 with $p > z$ is at most

$$X^{-1}\sum_{z<p\le\sqrt{X+b}}(X/p^2+1) \le \sum_{p>z}p^{-2} + O_b(X^{-1/2}).$$

Let $X \to \infty$ and then $z \to \infty$, keeping Q fixed. The product is positive because its factor at $p = 2$ is positive and the other deficits are summable.

For completeness, a bounded sequence with partial sum $A(x) = \delta x + o(x)$ has harmonic mean $\delta + o(1)$ on every $(U, V]$ with $\ell \to \infty$. Given $\epsilon > 0$, choose X_0 such that $|A(x) - \delta x| \le \epsilon x$ above X_0. Partial summation there costs $O(1 + \epsilon\ell)$ in the error, and the part below X_0 costs $O(\log(2X_0))$. Divide by $W = \ell + O(1)$ and then let $\epsilon \to 0$. Apply this to the joint indicator and to the periodic conditioning event, whose density is positive. This gives (30). □

In each regime giving (24), choose h large enough that the upper limiting conditional collision probability is below $c_{\mathrm{sf}}/2$. For this fixed h,

$$\begin{aligned}\mathbb{P}(A_k(dm) = 1,\ \mu^2(m+b) = 1 \mid R_h) \ge{}& \mathbb{P}(\mu^2(m+b) = 1 \mid R_h)\\ &- \mathbb{P}(A_k(dm) > 1 \mid R_h).\end{aligned}$$

Multiplying by $\mathbb{P}(R_h) \to \rho(h)$ gives

$$\liminf_{D,\ell\to\infty}\mathbb{P}(A_k(dm) = 1,\ \mu^2(m+b) = 1) \ge \rho(h)c_{\mathrm{sf}}/2 > 0.$$

For $U \ge (2D)^{C_r}$, the same conclusion follows from the fixed h in (26). This proves Theorem 1.2(3) for μ.

At $r = 0$, the collision probability is $o(1)$ without conditioning. The explicit test above and the triangle inequality show that the optimal correlation differs from $\mathbb{E}|f(m + b)|$ by

$o(1)$. Lemma 7.1 with $Q = 1$ gives $\mathbb{E}\mu^2(m+b) \to 1/\zeta(2) = 6/\pi^2$, completing the sparse-layer assertion.

At $r = 1/2$, for example,

$$J(r) = 0.954771\ldots < 1 < K(r) = 1.024504\ldots.$$

If $U = D$, a width $\ell = (\log D)^{0.98}$ already gives positive probability of recovering the whole cofactor. For the same layer and $\ell = \log\log D$, the exchange bound (23) loses too much at its boundary and the exponent in (28) is positive. Whether a positive proportion of products still determine the cofactor in this second example is open here. This identifies the remaining distinction between the sharp threshold for cancellation on all windows and the dependence of recovery on the location and width of an individual window.

AI assistance and finite checks. OpenAI's ChatGPT assisted in developing the cofactor and multiplicity arguments. OpenAI's Codex assisted in treating the sparse layers, the window dependence and the squarefree restriction, checking the input statements, and preparing the English exposition. ChatGPT also assisted with the bibliographic revision and proof exposition. Codex assisted with the results on additional observations, the reference checks and the source archive. The accompanying programs reproduce the threshold constants and verify finite allocation, squarefree, moment and decoding calculations.

References

[1] Idriss Olivier Bado. Average shift counts for the divisibility $\varphi(n) \mid n + a$. ResearchGate, 2026. 27 May 2026, Section 7, Proposition 5. DOI 10.13140/RG.2.2.29563.07202.

[2] Jean Bourgain, Peter Sarnak, and Tamar Ziegler. Disjointness of Moebius from horocycle flows. In *From Fourier Analysis and Number Theory to Radon Transforms and Geometry*, volume 28 of *Developments in Mathematics*, pages 67–83. Springer, 2013. DOI 10.1007/978-1-4614-4075-8_5. Theorem 2, p. 69.

[3] Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory*. Wiley-Interscience, second edition, 2006. DOI 10.1002/047174882X.

[4] Kevin Ford. The distribution of integers with a divisor in a given interval. *Ann. of Math. (2)*, 168:367–433, 2008. DOI 10.4007/annals.2008.168.367.

[5] Kevin Ford. Extremal properties of product sets. *Proc. Steklov Inst. Math.*, 303:220–226, 2018. DOI 10.1134/S0081543818080175.

[6] Kevin Ford. Joint Poisson distribution of prime factors in sets. *Math. Proc. Cambridge Philos. Soc.*, 173:189–200, 2022. DOI 10.1017/S0305004121000499. Lemma 2.1. author text.

[7] Kevin Ford, Mikhail R. Gabdullin, and Kyle Pratt. Integers with many divisors in a dyadic interval. arXiv:2611.15534, 2026. DOI 10.48550/arXiv.2611.15534.

[8] Imre Kátai. A remark on a theorem of H. Daboussi. *Acta Math. Hungar.*, 47:223–225, 1986. DOI 10.1007/BF01949145.

[9] Dimitris Koukoulopoulos. Localized factorizations of integers. *Proc. London Math. Soc. (3)*, 101(2):392–426, 2010. DOI 10.1112/plms/pdp056.

[10] Dimitris Koukoulopoulos. On the number of integers in a generalized multiplication table. *J. Reine Angew. Math.*, 689:33–99, 2014. DOI 10.1515/crelle-2012-0064. Theorem 1.7 and Section 2.1.

[11] Helmut Maier and Gérald Tenenbaum. On the set of divisors of an integer. *Invent. Math.*, 76:121–128, 1984. DOI 10.1007/BF01388495.

[12] Brice Pouly. An explicit bound below one half for a quartic Liouville correlation. Zenodo, 2026. DOI 10.5281/zenodo.24443184.

[13] Brice Pouly. Observations and Shift Decoding in Divisor Incidences. Zenodo, 2026. DOI 10.5281/zenodo.24578881.

[14] Brice Pouly. Parity information in a rough cofactor and exact reconstruction thresholds. Zenodo, 2026. DOI 10.5281/zenodo.24578849.

[15] Brice Pouly. Reconstruction and sign prediction from sampled prime factors. Zenodo, 2026. DOI 10.5281/zenodo.24443266. See also Part III, Sections 23–25.

[16] Jeremy Schlitt. Multiplication tables for integers with restricted prime factors. arXiv:2603.19212, 2026. Theorem 1.2 and Corollary 1.3. DOI 10.48550/arXiv.2603.19212.

[17] Atle Selberg. Note on a paper by L. G. Sathe. *J. Indian Math. Soc. (N.S.)*, 18:83–87, 1954. Theorem 4, pp. 85–87. Digitized original.

[18] Kannan Soundararajan and Max Wenqiang Xu. Central limit theorems for random multiplicative functions. *J. Anal. Math.*, 151:343–374, 2023. DOI 10.1007/s11854-023-0331-y.

[19] Terence Tao. The logarithmically averaged Chowla and Elliott conjectures for two-point correlations. *Forum Math. Pi*, 4:e8, 2016. DOI 10.1017/fmp.2016.6. Theorems 2 and 3, Corollary 5.

[20] Gérald Tenenbaum. Un problème de probabilité conditionnelle en arithmétique. *Acta Arith.*, 49:165–187, 1987. DOI 10.4064/aa-49-2-165-187.

[21] Gérald Tenenbaum. Crible d'ératosthène et modèle de Kubilius. In *Number Theory in Progress*, volume 2, pages 1099–1130. de Gruyter, 1999. DOI 10.1515/9783110285581.1099.

[22] Gérald Tenenbaum. *Introduction to Analytic and Probabilistic Number Theory*, volume 163 of *Graduate Studies in Mathematics*. American Mathematical Society, third edition, 2015. DOI 10.1090/gsm/163. Chapter II.6, Section 6.1, pp. 299–304.