hexagon:2611.01207v1

Visible prime factors under neighbouring critical divisor weights

Brice Pouly*

Abstract

A positive divisor weight marks one prime P of an integer $X < n \le 2X$. We compare the law of (n, P) before and after reweighting at $n + 2$. For fixed $t \ge 0$ and $u \to \infty$ with $u = o(\log X)$, the total variation distances for the exact cofactor n/P, the mark P and the origin n tend to 0, e^{-t} and 1, respectively. When further factors are revealed in uniform order, the distance equals the probability, under the first law, of exhausting the factorisation up to $o(1)$, uniformly over stopping rules. A second deletion compares all incomplete records. We also determine the joint law of n/P, the exact neighbouring mark and the neighbouring prime divisors up to $y = X^{o(1)}$. The approximation is in total variation, with conditional error tending to zero in mean over the cofactor. At $y = X^{1/V}$ the reference conditioned on a visible neighbouring mark has distance $e^{-u/V} + o(1)$. The mark, the other small primes and the full factor count give independent exponential, Poisson-process and Gaussian limits on their respective scales.

1 Which observations distinguish the two laws?

How much of an integer's factorisation must be observed to distinguish two populations whose weights differ at a neighbouring integer? The answer depends on which information is retained. An exact cofactor can have nearly the same law in both populations, while a sufficiently long list of exact prime factors reconstructs the origin and separates them. We quantify this change in total variation, which permits every test of the observed integer labels. All logarithms are natural.

For $m > 1$ and $s \ge 0$ put

$$H_s(m) = \prod_{p \mid m} (1 - p^{-s}), \qquad W_s(m) = \partial_s H_s(m) = \sum_{p \mid m} (\log p) p^{-s} \prod_{\substack{q \mid m \\ q \ne p}} (1 - q^{-s}).$$

Products and sums over p, q here run over distinct primes. Set $H_s(1) = 1$, $W_s(1) = 0$. The positive summands of W_s specify a choice of prime. At $s = 0$ the weight is the von Mangoldt function Λ. For $s > 0$, integers with several distinct prime factors carry weight as well.

Write $I = (X, 2X]$, $L = \log X$, $\mathbb{E}_X f = X^{-1} \sum_{n \in I} f(n)$ and fix $t \ge 0$. For each $p \mid n$ take the summand

$$w(n, p) = (\log p) p^{-t/L} \prod_{\substack{q \mid n \\ q \ne p}} (1 - q^{-t/L}).$$

On the pairs (n, P) give $\mathbb{P}_X$ mass proportional to $w(n, P)$. Give $\mathbb{Q}_X$ mass proportional to $w(n, P) W_{u/L}(n + 2)$. Each law is divided by its own total mass. The main regime is

$$X \to \infty, \qquad u = u(X) \to \infty, \qquad u = o(L). \tag{1}$$

*Research and writing with AI assistance.

We also give estimates with u fixed before taking X to infinity. The cofactor is the integer $D = n/P$, with one occurrence of P removed. We write $\mathbb{P}_X^O, \mathbb{Q}_X^O$ for the laws of an observation O and $d_{\mathrm{TV}}(\mu,\nu) = \frac{1}{2}\sum|\mu-\nu|$ for probabilities on discrete spaces. With equal prior probabilities, the best chance of identifying the population from O is $(1 + d_{\mathrm{TV}}(\mathbb{P}_X^O, \mathbb{Q}_X^O))/2$. For finite measures we use the full norm $\|\sigma\|_1 = \sum|\sigma|$.

For example, take $210 < n \le 420$, $t > 0$ and observe $D = 21$. The marked prime can be $11, 13, 17$ or 19. The first two choices give $n + 2 = 233$ and $275 = 5^2 \cdot 11$, respectively. Thus the neighbour can have a small prime factor for some of the integers producing the same observation. In a different experiment, start with the mark 13 and next reveal 7. Their product 91 is too small to be the origin. Revealing 3 as well gives 273, which lies in the sampling interval. For every squarefree origin, the product of revealed factors exceeds X precisely when all factors have been revealed. This elementary threshold will make reconstruction observable.

Theorem 1.1 (Exact factor observations)**.** *Under* (1)*, let O_j display P and up to $j-1$ further distinct prime divisors in uniformly random order without replacement. The list ends when the factors are exhausted. If N has the Poisson law of mean t, let C_N count the cycles of a uniform permutation of N letters, with $C_0 = 0$. Then*

Observation	$\lim_{X\to\infty} d_{\mathrm{TV}}(\mathbb{P}_X^O, \mathbb{Q}_X^O)$
full cofactor $D = n/P$	0
designated prime P	e^{-t}
ordered list O_j	$a_j(t) = \mathbb{P}(1 + C_N \le j)$
origin n	1

The error for O_j tends to zero uniformly in integers $j \ge 1$. More generally, after seeing P a rule may decide whether to request the next uniformly ordered factor, using its previous observations and independent randomness. For the observed stopped list O,

$$d_{\mathrm{TV}}(\mathbb{P}_X^O, \mathbb{Q}_X^O) = \mathbb{P}_X\left(\prod_{p \text{ displayed}} p > X\right) + o(1), \tag{2}$$

uniformly over these rules, including rules that depend on X. The rule controls whether revelation continues; each next factor is chosen uniformly from those not yet revealed.

At $t = 0$ almost all the first-site mass is on primes and the table gives $a_j(0) = 1$. For $t > 0$, revealing one more factor increases the limiting distance by

$$a_2(t) - a_1(t) = e^{-t}\sum_{k\ge1}\frac{t^k}{k!\,k}.$$

This is the one-cycle probability in the Poisson mixture. At $t = 1$ and equal priors, the optimal identification probabilities are approximately 0.684 for one observed factor and 0.926 for two. The origin's exact factorisation can therefore separate the laws, even though the whole integer left after deleting a single designated prime has asymptotically identical distributions.

The proof of the stopping assertion uses a second deletion: on a composite origin, hide one uniformly chosen unmarked prime and retain the original mark. This provides a comparison of the entire submeasures of incomplete observations. The elementary reconstruction threshold then identifies exactly where the two populations can separate.

The second question runs in the other direction. Having observed D, what is the law of the neighbour's factors? Write $s = u/L$. Under $\mathbb{Q}_X$ choose its mark $R \mid n + 2$ by the summands of $W_s(n + 2)$. Choose $y = X^{1/V}$ with $V = V(X) \to \infty$ and $V = o(u)$. Record

the set of distinct prime divisors of $n+2$ up to y and the mark if it is at most y, using a separate symbol $\star$ for a larger mark. Call this observation $\mathcal{F}_y$.

The reference law $\mathcal{M}_{s,y}$ on marked lists is as follows. Draw a mark R among the odd primes up to y with probabilities proportional to c_p. Include R in the set and each other odd prime up to y independently with probability q_p. The local parameters are defined for every odd prime by

$$c_p = \frac{(\log p)p^{-s}}{p-1-p^{-s}}, \qquad q_p = \frac{1-p^{-s}}{p-1-p^{-s}}. \tag{3}$$

Theorem 1.2 (The cofactor and the marked list)**.** *Under* (1)*, for every* $V \to \infty$ *with* $V = o(u)$ *and* $y = X^{1/V}$*,*

$$d_{\mathrm{TV}}\Big(\mathbb{Q}_X^{(D,\mathcal{F}_y)},\ \mathbb{P}_X^D \otimes \mathcal{M}_{s,y}\Big) \longrightarrow 0. \tag{4}$$

Consequently the conditional law of the marked list, given the exact value of D*, approaches* $\mathcal{M}_{s,y}$ *in mean total variation under* $\mathbb{Q}_X$*.*

At a lower cutoff the mark may lie outside the displayed list. To describe this event, put $C_s(z) = \sum_{2<p\le z} c_p$. For $s > 0$, the series $C_s(\infty)$ converges. Choose an odd prime R_s with mass $c_p/C_s(\infty)$ at p. Independently draw Bernoulli variables B_p with means q_p for the odd primes. Form the set

$$A_{s,y}^* = \{2 < p \le y : B_p = 1\} \ \cup\ (\{R_s\} \cap [2,y]), \qquad \widehat{\mathcal{M}}_{s,y} = \mathcal{L}(A_{s,y}^*, R_s).$$

The arithmetic set is $A_y = \{p \le y : p \mid n+2\}$. When $R > y$, the pair (A_y, R) still gives the exact prime, whereas $\mathcal{F}_y$ records $\star$.

Corollary 1.3 (Keeping the neighbouring mark at every cutoff)**.** *Assume* (1)*. For every deterministic* $Y = Y(X) \ge 3$ *with* $\log Y = o(\log X)$*,*

$$\sup_{2\le y\le Y} d_{\mathrm{TV}}(\mathbb{Q}_X^{(D,A_y,R)},\ \mathbb{P}_X^D \otimes \widehat{\mathcal{M}}_{s,y}) \longrightarrow 0.$$

The mean under $\mathbb{Q}_X^D$ *of the supremum of the conditional distances to* $\widehat{\mathcal{M}}_{s,y}$ *also tends to zero. For the observation* $\mathcal{F}_y$ *and the model whose mark is always visible,*

$$\sup_{3\le y\le Y} \Big| d_{\mathrm{TV}}(\mathbb{Q}_X^{(D,\mathcal{F}_y)},\ \mathbb{P}_X^D \otimes \mathcal{M}_{s,y}) - e^{-s\log y} \Big| \longrightarrow 0.$$

For example, at $y = X^{1/u}$ the cost of forcing a visible mark tends to e^{-1}. Keeping the exterior symbol yields a model for the entire list even at this cutoff. Section 5 proves the arithmetic comparison first at X^{1/V_0} with $V_0 \to \infty$, $V_0 \le \sqrt{u}$. For each exact cofactor m the remaining prime ranges over its own interval. Summing the absolute progression discrepancies over m gives the joint estimate, so tests may depend on both the row and the marked list. Projection then supplies all the smaller cutoffs. The conditional statements average over rows with their actual probabilities.

In particular, put $T = \log u$ and $Y = (\Omega(n+2) - T)/\sqrt{T}$, where Ω counts prime factors with multiplicity. We will prove

$$\mathbb{E}_{\mathbb{Q}_X} \sup_x \big| \mathbb{Q}_X(Y \le x \mid D) - \Phi(x) \big| \longrightarrow 0, \tag{5}$$

where Φ is the standard normal distribution function. At the scale of individual small factors the limit is different. If $z \ge 0$ and $P^-(m)$ denotes the least prime factor, then, in mean absolute difference,

$$\mathbb{Q}_X\big(s\log P^-(n+2) > z \mid D\big) \longrightarrow \exp\left(-z - \int_0^z \frac{1-e^{-v}}{v}\,\mathrm{d}v\right). \tag{6}$$

1.1 Related work and the information retained

Three observations occur in the literature: logarithmic factor sizes, exact small prime divisors, and exact integer cofactors. Billingsley's limit theorem and its proof by Arratia and Kochman concern logarithmic sizes for uniform integers [3, 2]; Granville develops their connection with Poisson statistics and cycles [14]. Arratia studies dependence between exact prime factors and couplings by insertions and deletions [1]. Haddad and Koukoulopoulos obtain more precise couplings for logarithmic factors [17]. The metric matters here: two distinct primes may have very close normalised logarithms while the point masses at their exact labels have total variation distance one.

The one-site law in Proposition 2.2 is an arithmetic realisation of classical partition laws. Its count is the number of cycles in a permutation of Poisson size, hence a Poisson mixture of the Ewens sampling formula with parameter one [9, 28]. Its full marked simplex admits a corresponding interpretation through sampling from a Poisson–Dirichlet partition, given after that proposition. Elboim and Gorodetsky study the relation between positive multiplicative weights and generalised Ewens measures [8]. For a comparison involving an exact cofactor, Sun's *Tilting Billingsley's model toward a giant prime: two phase transitions* [27, abstract] gives a total variation approximation for the cofactor of the *largest* prime under a weight favouring that prime. In the present experiment the derivative $W_{t/\log X}$ designates the prime to remove, and the reweighting acts on a second integer. At fixed t the number of first-site factors remains tight, as Section 2 shows.

For exact small factors, Ford gives a total variation approximation for the prime divisors of a shifted prime [10, Theorem 1, corrected version]. Gorodetsky's general Kubilius model for sieve-theoretic sequences [13, Theorem 2.3 and Corollaries 2.4–2.5] describes complete small-prime configurations from divisibility probabilities and their errors. His theorem permits a general arithmetic sampling law, including weighted laws; its use here requires uniform estimates for the marked weights and an absolute sum over cofactor rows. Our comparison retains the cofactor D jointly and introduces a neighbouring derivative weight with $u \to \infty$ arbitrarily slowly. Reweighting a total variation estimate by a weight bounded by $\log X$ can multiply its error by that factor when normalising masses are of constant order. An unweighted estimate alone therefore does not supply the asserted range. The proof in Section 5 uses relative sieve errors, sums absolute progression discrepancies over the cofactor rows, and recovers omitted configurations from an independently known total mass.

Reconstruction and sign prediction from sampled prime factors (DOI `10.5281/zenodo.24443266`) studies logarithmically weighted prime labels under harmonic sampling to recover an integer or predict its sign [26]. In Theorem 1.1 the target is the population from which the sample came. The remaining factors are revealed uniformly and a second deletion compares incomplete records. A different observation occurs in *Observations and Shift Decoding in Divisor Incidences* [25, Theorem 1.2]. There the sample is uniform on incidences (h, n, a, b) with $a \mid n$, $b \mid n + h$ in prescribed ranges. The observation $n + h$, optionally with b, is used to guess the hidden shift h. Here the shift is fixed at two and the prime labels carry the information.

Critical divisor weights and the localisation of parity defects [24, Theorem 1.1] studies the same pair of weights at $n, n + 2$. It locates the support of the defect in their scalar correlation with both smoothing parameters fixed on the critical scale. The present argument follows the cofactor n/P and takes u to infinity with X. Its progression estimates allow a different bounded test on every cofactor row. The shared sieve background is Bombieri's work on twin almost primes and the asymptotic sieve of Friedlander and Iwaniec [4, 11].

Sections 2–4 prove the one-site law and the deletion and reconstruction results. Sections 5–6 treat the neighbouring mark and its limiting laws. Appendix A specialises the progression theorem of Granville and Shao with the full factor count retained. Appendix B gives marked product laws for varying binary targets and for ternary targets, in the setting of the polynomial

Goldbach results of Perelli and Brüdern, Kawada and Wooley [22, 5].

2 The first integer and its marked factor

We need to discard small prime factors at the first site with a cost uniform in a moving second parameter. The two weights vanish at different rates on integers with a fixed small prime divisor.

Lemma 2.1 (Small factors and repeated factors)**.** *For fixed $0 < \eta < 1/10$, the contribution of $P^-(n) \le X^\eta$ to $\mathbb{E}_X W_{t/L}(n)$ is $O_t(\eta + 1/L)$. In $\mathbb{E}_X W_{t/L}(n) W_{u/L}(n+2)$ it is*

$$O_t(\eta + 1/L + ue^{-u/10}) \quad \textit{under } (1), \qquad O_{t,U}(\eta + 1/L) \quad (0 \le u \le U,\ U \textit{ fixed}).$$

After this removal, first-site squares cost $o_\eta(1)$ in either mean. All the bounds survive any multiplier of modulus at most one on the marked choices.

Proof. Take $K = 4(t+2)$, set $f(p^a) = \min(1, K \log p/L)$, and extend multiplicatively. Then $W_{t/L}(n) \ll_t Lf(n)$ for $n \in I$. In a marked term, each small unmarked prime supplies a factor at most t/K after division by its f-factor. Summing over k possible small marks costs $O_t(Lk(t/K)^{k-1})$. The logarithms of all large marks sum to $O(L)$. Also $0 \le W_s(m) \le \log m$ for every real $s \ge 0$.

Here is the upper sieve used below. Prescribe $p \mid n$ and $r \mid n+2$ with $p, r \le X^{1/10}$, or omit one or both conditions. In a compatible class the length is $V_0 \asymp X/[p,r]$, using 1 for any absent prime. Omit the prescribed primes and $2, 3$ from the sieve. For each remaining $\ell \le z = V_0^{1/80}$, retain the forbidden residues $0, -2$ independently with probabilities $1 - f(\ell)$ and ℓ^{-s}, respectively. The expected survival indicator bounds the product of the corresponding multiplicative factors from above, since factors for larger primes are at most one. Take the upper Selberg sieve with cutoff $V_0^{1/5}$ for each selection [12, Ch. 7]. If ν_ℓ residues are retained, its Euler density is $\prod(1 - \nu_\ell/\ell)$. In the unrestricted Selberg denominator the average $\log d$ is $\sum_{\ell \le z}(\nu_\ell/\ell)\log \ell \le 2\log z + O(1)$. Markov's inequality thus retains a fixed proportion of the denominator below the cutoff. The normalised remainder is $O(V_0^{-1} \sum_{d \le V_0^{2/5}} 6^{\omega(d)}) = O(V_0^{-1/2})$, by counting residue classes. The constants are uniform over the selections.

After averaging, the density factors are $1 - (1 - f(\ell) + \ell^{-s})/\ell$. Their product is $\ll_t s/L$ for $u \ge 1$, $s = u/L = o(1)$. Indeed $\sum_{\ell \le X} f(\ell)/\ell \ll_t 1$ by Mertens' estimates, and $\sum_\ell \ell^{-1-s} = \log(1/s) + O(1)$. The tail past z is bounded because $s \log z$ is bounded below. Removing a bounded number of primes costs a bounded factor. Writing $H_s^{(r)}$ for the product with the r-factor omitted, this proves

$$\mathbb{E}_X Lf(n) H_s(n+2) \ll_t s,$$
$$\mathbb{E}_X \mathbf{1}_{p \mid n,\ r \mid n+2} Lf(n) H_s^{(r)}(n+2) \ll_t sf(p)/(pr). \tag{7}$$

Only $p = r = 2$ can occur with equal prescribed primes, and its progression length changes the bound by at most a factor two. With no first-site prescription, the second bound holds with $f(p)/p$ omitted.

Neighbour marks $r > X^{1/10}$ contribute $O_t(ue^{-u/10})$, as follows by writing their sum as $H_s(n+2) \sum_{r \mid n+2,\ r > X^{1/10}} (\log r)/(r^s - 1)$ and applying the first bound. For the smaller marks multiply (7) by $(\log r) r^{-s}$ and sum. We have $\sum_r (\log r) r^{-1-s} \ll 1/s$ and $\sum_{p \le X^\eta} f(p)/p \ll_t \eta + 1/L$. This gives the moving-parameter estimate. For $0 \le u \le U$ majorise each W by $O_{t,U}(L)$ times the corresponding f, now with $K = 4(2 + t + U)$. Retaining each residue with probability $1 - f(\ell)$ gives a density $O_{t,U}(L^{-2})$, with the gain $O(f(p)/p)$ upon prescribing $p \mid n$. One forbidden residue gives $O_t(L^{-1})$ for the single weight. Lastly, the dyadic density

of n divisible by p^2 for some $p > X^\eta$ is $O(X^{-\eta} + X^{-1/2})$. The weights cost at most $O(L^2)$. Positivity proves the assertion about multipliers. □

Proposition 2.2 (The marked simplex)**.** *The mass $\mathbb{E}_X W_{t/L}(n)$ tends to 1. Under $\mathbb{P}_X$, $\Omega(n)$ converges in total variation to $1 + C_N$ of Theorem 1.1. On the stratum with j unmarked prime factors, list them in uniform order and put $x_i = \log q_i/L$ and $x_0 = \log P/L$. Their limiting* subprobability *for that stratum has density*

$$\frac{e^{-t}}{j!} \prod_{i=1}^{j} \frac{e^{tx_i} - 1}{x_i} \, \mathrm{d}x_1 \cdots \mathrm{d}x_j, \quad x_i > 0, \quad \sum_{i=1}^{j} x_i < 1, \quad x_0 = 1 - \sum_{i=1}^{j} x_i. \tag{8}$$

There is an atom e^{-t} at $j = 0, x_0 = 1$. Denote this marked factorisation law by $\mathcal{P}_t$.

Proof. First require squarefreeness and all factors to exceed X^η. The number of factors is now bounded in terms of η. For fixed, distinct $q_1, \dots, q_j$ the marked prime P is summed with weight $\log P$ in an interval of scale $X/(q_1 \cdots q_j) \gg X^\eta$, possibly cut at X^η. The prime number theorem gives an error $O_{\eta,A}((X/\prod q_i)L^{-A})$ uniformly in both endpoints. The cost of collisions is $o_\eta(1)$ by the square count. On the dyad $P^{-t/L} = e^{-t(1-\sum x_i)}(1 + O_t(1/L))$. Mertens' formula gives $\mathrm{d}x_i/x_i$ for the unmarked sums $1/q_i$ on fixed subintervals of $[\eta, 2]$. The factor $1/j!$ averages over orderings. Testing logarithmic rectangles, first on the truncated simplex and then letting $\eta \downarrow 0$ by Lemma 2.1, gives (8). The truncation faces may be approached from either side since they have zero limiting mass away from the stated atom.

For $j > 0$, expand the factors as power series with nonnegative coefficients. Terms with integers $m_i \geq 1$, $m = \sum_i m_i$, have integrated mass

$$e^{-t} \frac{t^m}{m!} \frac{1}{j!} \prod_{i=1}^{j} \frac{1}{m_i}, \qquad \int_{\sum x_i < 1} \prod_i x_i^{m_i - 1} \, \mathrm{d}\mathbf{x} = \frac{\prod_i (m_i - 1)!}{m!}.$$

The factor following the Poisson mass is the probability of the ordered list of cycle lengths $(m_1, \dots, m_j)$ in a uniform permutation of m letters, after ordering its cycles uniformly. This is the Ewens sampling formula at parameter one [9, 28]. Given the list, the log sizes have the Dirichlet law with parameters $(m_1, \dots, m_j, 1)$, with the last coordinate assigned to P. Summing gives mass one. Lemma 2.1 makes the count tight, and the point probabilities on each fixed stratum give the stated total variation convergence, including proper powers. □

For $0 \leq z < 1$, the limiting probability that $x_0 > z$ is e^{-tz}. Indeed restricting the simplex to $\sum x_i < 1 - z$ multiplies each integrated term of total degree m by $(1 - z)^m$. This logarithmic-size law has a continuous part and an endpoint atom. The mark in Theorem 1.1 is observed as an exact prime.

Sampling a classical random partition. There is a useful probabilistic construction of the whole law $\mathcal{P}_t$. A Poisson–Dirichlet partition with parameter one, denoted PD(1), is obtained by ranking the masses

$$U_i \prod_{h<i} (1 - U_h), \qquad i \geq 1,$$

where the U_i are independent uniform variables on $(0, 1)$. Independently draw $N \sim \mathrm{Poisson}(t)$ and sample N points from this partition. Retain the masses of the occupied blocks in uniform order, and collect the mass of all unoccupied blocks into a distinguished residual coordinate.

Conditional on $N = m$, the occupied sample counts $(m_1, \dots, m_j)$ have the ordered cycle probabilities used above. Conditional on these counts, the occupied masses and the residual mass have distribution Dirichlet$(m_1, \dots, m_j, 1)$; this follows from the predictive form of the

Ewens sampling formula [28, Section 2.4, equations (6)–(7)]. Multiplying these probabilities and Dirichlet densities by the Poisson mass, then summing over the m_i, gives exactly (8). Thus the residual coordinate represents a total unobserved mass in this construction. Proposition 2.2 identifies that classical law as the limit of a designated arithmetic prime and its complementary factors under $W_{t/\log X}$.

3 Deleting one prime factor

Put $S = 2\prod_{p>2}(1-(p-1)^{-2})$, the singular series for the pair $n, n+2$. For the cofactor define the unnormalised measures

$$\mu = \mathbb{E}_X \sum_{p\mid n} w(n,p)\delta_{n/p}, \qquad \nu_u = \mathbb{E}_X \sum_{p\mid n} w(n,p)W_{u/L}(n+2)\delta_{n/p}.$$

On summing over all cofactors, $\nu_u(\mathbf{1})/S = \mathbb{E}_X[W_{t/L}(n)W_{u/L}(n+2)]/S$ is precisely the scalar correlation in *Critical divisor weights and the localisation of parity defects* [24, equation (2)]. Testing against a function of n/p retains its distribution among the integer cofactors.

We shall also need a second deletion. On squarefree composites, retain the original mark P, choose q uniformly among the other prime factors, and record $(n/q, P)$. Use the same notation for these two submeasures when this observation is specified.

Proposition 3.1 (One free prime)**.** *For either deletion, at each fixed $u \geq 1$ and $\alpha = 4/\pi$,*

$$\limsup_{X\to\infty} \|\nu_u - S\mu\|_1 \ll_t (1+u)^{\alpha} e^{-u/2}, \tag{9}$$

$$\limsup_{X\to\infty} \left\| \int_0^u (\nu_v - S\mu)\,\mathrm{d}v \right\|_1 \ll_t (1+u)^{\alpha} e^{-u/2}. \tag{10}$$

Along every sequence (1), $\|\nu_u - S\mu\|_1 \to 0$. *All three assertions hold after adjoining the index of any prespecified partition of I into at most L^{B_0} intervals, for fixed B_0. Their endpoints may depend on X.*

In particular the two-site mass tends to S in the moving regime, and the cofactor laws can be normalised in (9) with the same order of error at all sufficiently large fixed u. With $D = m$ and equal-length interval labels, the first deletion localises P to relative width $O(L^{-B_0})$ for $B_0 > 0$, while the enlarged observation still has distance tending to zero.

3.1 Distribution in progressions

Test either measure against an arbitrary function b of modulus at most one, which may depend on X. Let $a_\eta(n)$ denote its total first-site weight on squarefree n with $P^-(n) > X^\eta$. For the first observation a_η agrees there with the convolution of the prime sequence $(\log p)p^{-t/L}\mathbf{1}_{p>X^\eta}$ and the sequence

$$b(m)H_{t/L}(m)\mathbf{1}_{\mu_{\mathrm{M}}^2(m)=1,\ P^-(m)>X^\eta}.$$

Here μ_{M} is the Möbius function and $P^-(1) = \infty$. In the second case use the prime sequence $(1-q^{-t/L})\mathbf{1}_{q>X^\eta}$ and, on squarefree rough $m > 1$, the sequence

$$\frac{1}{\Omega(m)} \sum_{p\mid m} w(m,p)b(m,p).$$

In fact $w(mq,p) = (1-q^{-t/L})w(m,p)$ for $q \nmid m$ and $\Omega(m)$ is the number of choices for the deleted unmarked prime. Call either convolution a^*. The two sides can differ only at a square of a prime $> X^\eta$. Both are $O_\eta(L)$ because the number of possible prime factors is

bounded on the rough set. Against the final expanded divisor test $L\sum_{d|n+2}\mu_{\mathrm{M}}(d)\phi(\log d/L)$ the difference costs $o_\eta(1)\|\phi\|_\infty$ in the dyadic average, since $\tau(n+2)\ll_\eta X^{\eta/2}$.

For every fixed $0<\theta<1/2$, $A,J>0$, the convolution satisfies

$$\sum_{d\le X^\theta}\tau(d)^J\max_{(a,d)=1}\left|\sum_{\substack{n\in I\\ n\equiv a\ (d)}}a^*(n)-\frac{1}{\varphi(d)}\sum_{n\in I}a^*(n)\right|\ll_{A,J,t,\eta,\theta}XL^{-A}. \tag{11}$$

The term $m=1$ uses the maximal Bombieri–Vinogradov theorem and partial summation in the prime. For the remaining terms, divide the two variables into boxes of relative width L^{-K} and count the boxes strictly inside the product interval. The boundary products lie in two intervals of total length $O(XL^{-K})$. Their summed cost over moduli is $O_\eta(L^{O(1)}(XL^{-K}+X^\theta))$, by counting integers in a class and in the principal average. Each factor of an interior box is supported in a dyadic range of size at least a fixed power of X, and the two sizes have product comparable to X. One factor is a prime sequence with a smooth monotone weight or its logarithmic multiple. Partial summation then gives Siegel–Walfisz with arbitrary logarithmic savings. Removing primes dividing an auxiliary modulus r costs $O(L\omega(r))$, admissible in the divisor-weighted Siegel–Walfisz hypothesis. The other coefficients are arbitrary of size $O_\eta(L)$. The convolution theorem of Motohashi [20], as formulated in *Variants of the Selberg sieve, and bounded intervals containing many primes* [23, Claim 2.6 and Theorem 2.8 in the arXiv text], now gives arbitrary logarithmic savings summed over d. Choose the saving after K to absorb all the boxes. To pass to the total mass in the main term, observe that for each prime $\ell>X^\eta$ dividing n, $\sum_{d\le X^\theta,\,\ell\mid d}1/\varphi(d)\ll L/(\ell-1)$. This costs $O_\eta(X^{1-\eta}L^{O(1)})$. To introduce τ^J, apply Cauchy–Schwarz between the unweighted discrepancy and the discrepancy with τ^{2J} estimated by the trivial progression bound $O_\eta(L(X/d+1+X/\varphi(d)))$. The latter sums to $XL^{O_J(1)}$.

For the extra interval label, perform the argument on each labelled interval with its own test. The boundary cost is uniform in its endpoints, and the $m=1$ term uses the maximal version at both endpoints. Choose the precision in (11) larger by B_0+2 before summing. Since the intervals are disjoint, absolute bounds on the exact coefficients and square exceptions are paid for just once. Henceforth a_η can include these labels.

3.2 The divisor coordinate at the second site

For a smooth test ϕ put

$$T_X(\phi)=L\mathbb{E}_X a_\eta(n)\sum_{d\mid n+2}\mu_{\mathrm{M}}(d)\phi(\log d/L),\qquad \bar a_\eta=\mathbb{E}_X a_\eta.$$

Its support is contained in $[0,1+O(1/L)]$. If $\widehat{T}_X(\xi)=T_X(e^{-\mathrm{i}\xi x})$, then

$$|\widehat{T}_X(\xi)|\ll_t\exp\left(\sum_{p\le 3X}|1-e^{-\mathrm{i}\xi\log p/L}|/p\right)\ll_t(1+|\xi|)^2. \tag{12}$$

Indeed $|a_\eta|\le W_{t/L}\ll_t Lf(n)$ and the neighbour's Fourier multiplier has modulus $\prod_{p\mid n+2}|1-e^{-\mathrm{i}\xi\log p/L}|$. Nair and Tenenbaum's *Short sums of certain arithmetic functions* [21, Corollary 3] applies to $n,n+2$. The forms exclude two residues at odd primes, giving density $O(L^{-2})$, and the two multiplicative majorants are bounded uniformly by $2^{\omega(n)}\ll_\epsilon n^\epsilon$. The first prime sum is $O_t(1)$ and the higher prime powers cost $O(1)$. For the last inequality in (12) bound each numerator by $\min(2,|\xi|\log p/L)$, splitting at $\exp(L/(1+|\xi|))$ and using Mertens' formula. For each fixed frequency, the exponent tends to

$$\int_0^{|\xi|}\frac{2|\sin(r/2)|}{r}\,\mathrm{d}r=\alpha\log(1+|\xi|)+O(1).$$

For the first assertion omit $\log p/L < \epsilon$ at cost $O(|\xi|\epsilon)$ and use Mertens on the remaining compact interval. For the second, the centred periodic numerator has a bounded primitive.

Fourier inversion in (12) bounds T_X by a common seminorm involving four derivatives of tests on a fixed compact neighbourhood of its support. From any sequence we may extract $T_X \to T$ as distributions and $\bar{a}_\eta \to b_0$. The limit satisfies $|\widehat{T}(\xi)| \ll_t (1+|\xi|)^\alpha$. To locate it, take a test supported strictly below $1/2$ on the nonnegative axis, choose a fixed $\theta < 1/2$ above that support, and use (11). The class is -2 for odd d, and $n+2$ is odd on these rows. The model is

$$\bar{a}_\eta L \sum_d c(d)\phi(\log d/L), \quad c(d) = \mu_{\mathrm{M}}(d)\mathbf{1}_{2\nmid d}/\varphi(d).$$

The replacement error is $o_\eta(1)\|\phi\|_\infty$ with arbitrary logarithmic saving from (11) and a power saving for square collisions. Replacing the total mass of a^* by that of a_η also costs a power saving, since $\sum_{d\le X} 1/\varphi(d) \ll L$.

For clarity, $F(z) = \sum c(d)d^{-z} = K(z)/\zeta(1+z)$, where the Euler product for F is $\prod_{p>2}(1-p^{-z}/(p-1))$ and $K(0) = S$. Away from two the factor of K is $1 - p^{-1-z}/((p-1)(1-p^{-1-z}))$, so its coefficients are absolutely summable with weight $d^{1/4}$. The classical prime number theorem for $\mu_{\mathrm{M}}(m)/m$ and convolution give $C(v) = \sum_{d\le e^v} c(d) = O_A((1+v)^{-A})$ for every A. As $F(z) = z\int_0^\infty C(v)e^{-zv}\,\mathrm{d}v$, we have $\int C(v)\,\mathrm{d}v = F'(0) = S$. Partial summation yields

$$L\sum c(d)\phi(\log d/L) = -\int_0^\infty C(v)\phi'(v/L)\,\mathrm{d}v \longrightarrow -S\phi'(0).$$

Consequently $T = Sb_0\delta_0' + D_1$, with D_1 supported in $[1/2,1]$ and $|\widehat{D}_1(\xi)| \ll_t (1+|\xi|)^\alpha$.

Completion of the proof of Proposition 3.1. The tests $-xe^{-ux}$ and e^{-ux} give $\mathbb{E}_X a_\eta W_{u/L}(n+2)$ and $L\mathbb{E}_X a_\eta H_{u/L}(n+2)$, respectively. The latter is the primitive from 0 to u of the former. On a neighbourhood of $\operatorname{supp} D_1$, write $e^{-ux} = e^{-u/2}h(u(x-1/2))$ with h Schwartz and $h(v) = e^{-v}$ for $v \ge -1/2$. Fourier inversion after dilation costs $O_t((1+u)^\alpha e^{-u/2})$ for $u \ge 1$. Multiplication by x replaces $h(v)$ by $(1/2 + v/u)h(v)$ and satisfies the same bound. The origin distribution supplies Sb_0 and Sb_0u for these two tests. Remove η after $X \to \infty$ using Lemma 2.1, uniformly on $[0,u]$ for the primitive at fixed u. The Fourier constants are independent of η and of b.

For moving u, we return to the progression estimate (11) before extracting a distributional limit. This gives a separate argument for $u = u(X)$ with no interchange of the fixed-u limits. Take $0 < \gamma < \theta < 1/2$ and a smooth cutoff equal to one near $[0,\gamma]$ with support below θ. The complement of $-xe^{-ux}$, localised to a fixed neighbourhood of the support, has four-derivative bound $O((1+u)^4 e^{-\gamma u})$. On the cutoff, (11) applies before extraction, uniformly for bounded tests. Restoring the complementary part of the model costs $o(1)$ by partial summation with $C(v)$ beyond γL. Indeed the x-test there has uniformly bounded variation for $u \ge 1$, so the factor L is absorbed by the logarithmic saving. The full model has value $\bar{a}_\eta F'(s) = \bar{a}_\eta(S + O(s))$. This proves $\mathbb{E}_X a_\eta W_s(n+2) = S\bar{a}_\eta + o_\eta(1)$. Lemma 2.1 now permits $\eta \downarrow 0$ along (1). In all three assertions the test may be chosen at each X to attain the ℓ^1 norm. This proves the proposition. □

4 Reconstruction and stopping after an observed factor

The separating event uses only divisibility by a finite set of primes. For $z \ge 3$ set $d_z = \prod_{2<\ell\le z}\ell$ over primes and $E_z = \{(n+2, d_z) > 1\}$. In either (1) or the regime of fixed $u > 0$,

$$\lim_{X\to\infty} \mathbb{P}_X(E_z^c) = \prod_{2<\ell\le z}\left(1 - \frac{1}{\ell-1}\right) =: \alpha_z, \qquad \lim_{X\to\infty} \mathbb{Q}_X(E_z) = 0. \tag{13}$$

The first-site weight is asymptotically equidistributed on the reduced classes modulo any fixed d_z by applying (11) for the constant cofactor test, and Lemma 2.1 removes the cutoff. The same proof works for the submeasure $\Omega(n) \le j$, by testing $1 + \Omega(D) \le j$. Among the reduced classes at ℓ, exactly one is -2. The product α_z tends to zero.

For the second assertion one can either use the small-factor model proved in Section 5 in the moving case, or argue directly. For each fixed odd ℓ, the mark ℓ at $n+2$ contributes $O_{t,\ell}(s)$ by the no-p form of (7). For another mark $r \le X^{1/10}$, $r \ne \ell$, extract $1-\ell^{-s} = O_\ell(s)$ and apply the same sieve with $\ell r \mid n+2$. The remaining contribution is $O_{t,\ell}(s/r)$ before multiplying by $(\log r) r^{-s}(1-\ell^{-s})$. Summing in r gives $O_{t,\ell}(s)$. All larger marks cost $o(1)$. For fixed $u > 0$ the fixed-parameter version of the sieve gains $O_{t,u,\ell}(L^{-1})$ by forcing $\ell \mid n+2$ in the majorant for $W_{u/L}$. The denominator stays bounded below in this case. The dyadic Chen-prime theorem in Green and Tao's *Restriction theory of the Selberg sieve, with applications* [16, Theorem 6.1] supplies $\gg X/L^2$ primes n such that $n+2$ is prime or a product of two primes each $> X^{1/11}$. Both weights are $\gg_{t,u} L$ on this population (a repeated prime also suffices). For moving u the total mass is $S + o(1)$ by Proposition 3.1. This proves (13). In particular the full-origin distance tends to one, in both regimes, by taking X and then z to infinity.

Proof of Theorem 1.1. The cofactor assertion follows upon normalising Proposition 3.1 with Proposition 2.2. Both measures give $o(1)$ to nonsquarefree origins. To see this, discard factors $\le X^\eta$, use the square count, then take $\eta \downarrow 0$. We can work on squarefree origins up to this common error.

Generate a uniform ordering of the unmarked factors, whether or not the rule eventually asks for them. If the origin is composite, choose the last unmarked factor q first. Conditional on (n, P, q), the preceding unmarked factors form a uniform ordering of the prime divisors of $m = n/q$ other than P. Thus the following simulation uses just (m, P) and independent randomness. Display P, then follow the rule using the factors of m in uniform order. If the rule stops before asking for the hidden last factor, return the stopped transcript. Otherwise discard this simulation outcome. This construction gives exactly the submeasure of incomplete transcripts because the rule decides whether to request the last factor before learning its value. It also covers the case where the rule stops immediately after seeing P. On prime origins that mark already completes the record. Formally the simulation is a common subprobability kernel, so the ℓ^1 distance between its two output submeasures is bounded by the ℓ^1 distance between its inputs. The second deletion of Proposition 3.1, with the two full normalising masses, makes the latter $o(1)$ uniformly in the rule.

Let G be the set of transcripts whose product exceeds X. A proper subproduct of a squarefree $n \in I$ is at most $n/2 \le X$, so G identifies the completed transcripts from their values. The overlap of the two submeasures on G, namely the sum of their pointwise minima, is at most

$$\mathbb{P}_X(E_z^c) + \mathbb{Q}_X(E_z) + o(1).$$

Indeed the product on G is the actual origin, and we split the transcripts by the event E_z applied to that product. By (13) this overlap tends to zero as X then z tend to infinity, uniformly over the rules. The overlap on G^c is $\mathbb{P}_X^Q(G^c) + o(1)$ by the subprobability kernel just proved. Subtract the total overlap from one to get (2). An independent random seed may also be retained in the transcript. Condition on its value and integrate the same uniform bound.

For a deterministic budget j, $\mathbb{P}_X^Q(G) = \mathbb{P}_X(\Omega(n) \le j) + o(1)$, uniformly in j. Proposition 2.2 gives uniform convergence of these probabilities to $a_j(t)$. At $j = 1$ only the empty permutation of Poisson size contributes. At $j = 2$ a nonempty permutation of k letters contributes with probability $1/k$. This proves the remaining assertions. □

For every fixed $u > 0$ and $j \geq 1$ the completed-transcript test yields the bound

$$\liminf_{X\to\infty} d_{\mathrm{TV}}(\mathbb{P}_X^{O_j}, \mathbb{Q}_X^{O_j}) \geq a_j(t).$$

Indeed $\mathbb{P}_X(G) \to a_j$ and the probability of $G \cap E_z^c$ is bounded above by $\mathbb{P}_X(E_z^c)$. For all sufficiently large fixed u the normalised transfer estimate also gives

$$\limsup_X \sup_{j\geq 1} |d_{\mathrm{TV}}(\mathbb{P}_X^{O_j}, \mathbb{Q}_X^{O_j}) - a_j(t)| \ll_t (1+u)^{4/\pi} e^{-u/2}.$$

For this last bound the common incomplete-transcript kernel has error of the indicated order, and the completed overlap still tends to zero. An origin interval label as in Proposition 3.1 may be supplied before the rule begins and included in the output. The proof of (2) then applies with the enlarged input $(m, P, \text{interval label})$.

5 The neighbour's marked list with the cofactor retained

The arithmetic part of the proof gives Theorem 1.2 at cutoffs with $V \to \infty$ and $V \leq \sqrt{u}$. We sum the progression errors with the cofactor fixed and use an independent total-mass estimate to recover configurations omitted by the sieve. Controlling the tail through its model mass is part of the classical Kubilius method [13, Section 2.2]. We will then retain the exact mark and project the prime list to prove Corollary 1.3 and the full range of Theorem 1.2.

5.1 The prime progression estimate in each row

Fix $\eta \in (0, 1/10)$, discard nonsquarefree first sites and primes at most X^η there, and write $n = mp$ with $p = P$. The row label is $D = m$, where m is squarefree, $P^-(m) > X^\eta$, $m \leq 2X^{1-\eta}$, including $m = 1$. Temporarily allow $p \mid m$. There are at most $O_\eta(1)$ such primes in each row. The extra total contribution is $O_\eta(X^{-\eta}L^2)$ even after inserting the second weight. For each row set

$$I_m = (X/m, 2X/m] \cap (X^\eta, \infty), \quad A_m = \frac{H_{t/L}(m)}{X} \sum_{p\in I_m} (\log p) p^{-t/L}, \quad A = \sum_m A_m.$$

Sums indexed by p here and below are over primes. The A_m are one-site masses, with no lower bound imposed on individual rows. For $Q \leq X^\theta$, $\theta = \eta/8$, every prime divisor of Q is below X^η, so $(m, Q) = 1$. For a unit class a let $r_m(Q, a)$ be the row sum restricted to $mp \equiv a \bmod Q$ minus $A_m/\varphi(Q)$. For any fixed $J, B > 0$,

$$\sum_m \sum_{Q\leq X^\theta} \tau(Q)^J \max_{(a,Q)=1} |r_m(Q, a)| \ll_{\eta,t,J,B} L^{-B}. \tag{14}$$

Indeed $Z = X/m \geq X^\eta/2$ and X^θ lies below $\sqrt{Z}/(\log Z)^C$ by a power of X, for any fixed C. On the prime variable use the maximal Bombieri–Vinogradov theorem on intervals up to $2Z$ and partial summation of $p^{-t/L}$. The interval cut at X^η causes no loss. All eligible primes exceed Q, and the residue is $a\bar{m}$. The same bound follows from the usual Λ version by removing higher powers, using $X^\theta Z^{1/2} L^{O(1)} \ll ZL^{-B'}$ in each row. The principal mass can use the actual prime count with the same weight. To insert τ^J, interpolate by Cauchy–Schwarz against the trivial bound as in (11). Thus the unnormalised error for each row is $O_{\eta,t,J,B'}(ZL^{-B'})$. Multiply by $H_{t/L}(m)/X \leq 1/X$ and sum $1/m \ll L$ to prove (14). This order of summation allows a different bounded test on every pair consisting of a row and a marked list.

5.2 The local probability measure

Start with independent odd-prime divisibility indicators of probabilities $1/(p-1)$ for $p \le y$. Weight each set of present primes by the individual marked summands of the divisor weight restricted to that set. Its total mass is

$$Z'_y(s) = Z_y(s) \sum_{2<p\le y} c_p, \qquad Z_y(s) = \prod_{2<p\le y} \left(1 - \frac{p^{-s}}{p-1}\right).$$

For a prescribed mark r, normalising each other prime gives exactly the inclusion probability q_p of (3). The relative weights of the possible marks are c_r, proving the construction of $\mathcal{M}_{s,y}$.

In our regime $Z'_y(s) = S + o(1)$ and $\sum_{2<p\le y} c_p \sim 1/s$. For the second assertion,

$$c_p = (\log p)p^{-1-s} + O((\log p)/p^2).$$

Mertens' estimate $\sum_{2<p\le z}(\log p)/p = \log z + O(1)$ and partial summation give, uniformly for $z \ge 2$ and $0 < s \le 1$,

$$C_s(z) = \frac{1-z^{-s}}{s} + O(1), \qquad C_s(\infty) = \frac{1}{s} + O(1), \qquad 1 - \frac{C_s(z)}{C_s(\infty)} = z^{-s} + O(s) \quad (s \to 0).$$

Indeed the bounded error in Mertens' estimate contributes $O(z^{-s} + s\int_1^z v^{-1-s}\,\mathrm{d}v) = O(1)$. The local correction in c_p is absolutely summable. At the arithmetic cutoff in this proof, $s\log y = u/V \to \infty$. For the first assertion extend the product to all odd primes. By Section 3 its value is $F(s) \sim Ss$. The omitted logarithmic tail is $O(\sum_{p>y} p^{-1-s}) = o(1)$, by the preceding weighted estimate or partial summation and $s\log y \to \infty$. This proves the mass assertion.

5.3 Summing marked configurations

For a set of included odd primes up to y let d be their squarefree product, and let $r \mid d$ be its designated mark. Under $\mathcal{M}_{s,y}$ we have $\mathbb{E}\log d \ll \log y$, because the mark adds at most $\log y$, and $q_p \ll 1/p$ for the others. Call $d \le X^{\theta/4}$ good. The model mass of the other lists is $O_\eta(1/V)$.

For a good pair (d, r) impose $mp \equiv -2 \bmod d$ and exclude $mp \equiv -2 \bmod \ell$ at all remaining odd primes $\ell \le y$ outside d. In row m the main term for the unweighted list is

$$A_m \, \frac{1}{\varphi(d)} \prod_{\substack{2<\ell\le y \\ \ell\mid d}} \left(1 - \frac{1}{\ell - 1}\right).$$

The fundamental lemma of the dimension-one sieve, with level $D_0 = X^{\theta/4}$ on the excluded primes, gives upper and lower bounds with relative error $o_\eta(1)$ in this main term plus summed progression remainders [12, Ch. 11]. Indeed $\log D_0/\log y = \theta V/4 \to \infty$, and the dimension bound from Mertens is uniform when primes of d are deleted. The remainder for a sieve modulus e uses $r_m(de, -2)$. Each coefficient of the fundamental-lemma weights is bounded in modulus by one.

Multiply by the nonnegative truncated mark weight

$$(\log r)r^{-s} \prod_{\substack{\ell\mid d \\ \ell\ne r}} (1 - \ell^{-s}) \le O(L).$$

For a fixed product de the choices of d, e, r number at most $\tau(de)^{O(1)}$. Since $de \le X^{\theta/2}$, (14) bounds the sum of the absolute remainders over all *rows and configurations* by $o_\eta(1)$:

choose J to absorb this divisor power and then B to absorb the factor $O(L)$ from the mark. The relative sieve errors sum to $o_\eta(1)$ times $AZ'_y(s) = O(1)$ by positivity. For the actual second-site mark $r \le y$, the factors $> y$ multiply this truncated weight by a number in $[1 - O(Ve^{-u/V}), 1]$ because there are at most $2V$ such factors and each omitted $\ell^{-s} \le e^{-u/V}$. The loss is $o(1)$ since $V \le \sqrt{u}$. Thus the joint *unnormalised* mass on the good lists is within $o_\eta(1)$ in full ℓ^1 of the product of $(A_m)_m$ and $Z'_y(s)\mathcal{M}_{s,y}$ restricted to these lists.

The total arithmetic mass on the truncated first-site population is $SA + o_\eta(1)$ by the moving-parameter proof of Proposition 3.1, before removing η. This mass estimate is independent of the marked-list approximation just obtained. The model has that same total mass up to $o(1)$. Their masses agree on the good lists and the omitted model mass tends to zero. Positivity therefore makes the remaining arithmetic mass, including cases with mark $> y$, $o_\eta(1)$. The arguments cover the case $t = 0$, where effectively there is just the row $m = 1$. Finally remove η by Lemma 2.1, restore actual one-site cofactor masses, and divide by the two normalising masses. This proves (4) for $V \le \sqrt{u}$.

5.4 The exact mark and its projection to a smaller list

Fix Y as in Corollary 1.3. Work first at the larger cutoff

$$Z = \max\{Y, X^{1/\sqrt{u}}\} = X^{1/V_0}, \qquad V_0 = \min\{\log X/\log Y, \sqrt{u}\}.$$

Here $V_0 \to \infty$ and $V_0 \le \sqrt{u}$. Let $\epsilon_{X,Z}$ be the joint distance in (4) at Z. The preceding sieve argument gives $\epsilon_{X,Z} \to 0$. Its reference law $\mathcal{M}_{s,Z}$ is supported on visible marks. Giving every exterior arithmetic mark its exact prime label preserves the distance to this reference: on visible marks the masses agree with their former values, and on the exterior symbols the reference has zero mass. Splitting the latter atoms merely distributes their arithmetic mass among the exact labels.

In the global model $(A^*_{s,Z}, R_s)$, conditioning on $R_s \le Z$ gives exactly $\mathcal{M}_{s,Z}$. The complementary component is disjoint and has mass $\delta_s(Z) = 1 - C_s(Z)/C_s(\infty) = Z^{-s} + O(s) = o(1)$, since $s \log Z \ge \sqrt{u}$. Therefore

$$\eta_{X,Z} := d_{\mathrm{TV}}(\mathbb{Q}_X^{(D,A_Z,R)},\ \mathbb{P}_X^D \otimes \widehat{\mathcal{M}}_{s,Z}) \le \epsilon_{X,Z} + \delta_s(Z) = o(1).$$

Intersecting the list with $[2, y]$ and keeping R gives the corresponding laws at every $y \le Y$. All these distances are bounded by the same $\eta_{X,Z}$.

The same projection acts on conditional laws at each D. Write $K_Z(D) = \mathcal{L}_{\mathbb{Q}_X}(A_Z, R \mid D)$. Disintegration and marginalisation give

$$\begin{aligned}\mathbb{E}_{\mathbb{Q}_X} \sup_{2\le y\le Y} d_{\mathrm{TV}}(\mathcal{L}_{\mathbb{Q}_X}(A_y, R \mid D), \widehat{\mathcal{M}}_{s,y}) &\le \mathbb{E}_{\mathbb{Q}_X} d_{\mathrm{TV}}(K_Z(D), \widehat{\mathcal{M}}_{s,Z}) \\ &= d_{\mathrm{TV}}(\mathbb{Q}_X^{(D,A_Z,R)}, \mathbb{Q}_X^D \otimes \widehat{\mathcal{M}}_{s,Z}) \\ &\le \eta_{X,Z} + d_{\mathrm{TV}}(\mathbb{Q}_X^D, \mathbb{P}_X^D) \le 2\eta_{X,Z}.\end{aligned}$$

At each X only finitely many different lists occur as y varies. No lower bound on any individual cofactor probability enters this estimate.

It remains to compute what censoring does to the reference. Let $\mathcal{B}_{s,y}$ be the law of the independent Bernoulli list on odd primes up to y. For $y \ge 3$ the censored law in the global model is

$$(1 - \delta_s(y))\mathcal{M}_{s,y} + \delta_s(y)(\mathcal{B}_{s,y} \otimes \delta_\star).$$

Indeed, if $R_s > y$, all visible Bernoulli variables keep their original law. The mark distinguishes the two components, so their mixture is at distance exactly $\delta_s(y)$ from $\mathcal{M}_{s,y}$. The reverse triangle inequality, applied after the same censoring in both experiments, gives

$$\left| d_{\mathrm{TV}}(\mathbb{Q}_X^{(D,\mathcal{F}_y)}, \mathbb{P}_X^D \otimes \mathcal{M}_{s,y}) - \delta_s(y) \right| \le \eta_{X,Z}.$$

The uniform estimate $\delta_s(y) = e^{-s\log y} + O(s)$ completes Corollary 1.3. At $y = X^{1/V}$ the exponent is $s\log y = u/V$. Thus $V = o(u)$ recovers Theorem 1.2, including its conditional assertion by disintegration. When $u/V \to c < \infty$, the full distance to the always-visible model tends to e^{-c}. The error in the profile is additive. On the exterior event, whose mass then tends to $e^{-c} > 0$, one can also condition:

$$d_{\mathrm{TV}}(\mathcal{L}_{\mathbb{Q}_X}(D, A_y \mid R > y),\ \mathbb{P}_X^D \otimes \mathcal{B}_{s,y}) \longrightarrow 0.$$

Indeed, restricting two laws at distance η to a common event of probability at least b under both and normalising gives distance at most $2\eta/b$. The product law here describes the prime divisors one actually sees when the mark is outside the window.

Multiplicity. The probability of a squared prime factor at the neighbour tends to zero under $\mathbb{Q}_X$. For this separate estimate choose a cutoff satisfying $V \to \infty$, $V \le \sqrt{u}$; it need not be the cutoff used in Theorem 1.2. On retained good lists, sum over $\ell \mid d$ after imposing the extra modulus ℓ (thus $mp \equiv -2 \bmod d\ell$). The upper sieve just used has principal term multiplied by $\varphi(d)/\varphi(d\ell) = 1/\ell$. Since $d, e \le X^{\theta/4}$ and $\ell \le y = X^{o(1)}$, the moduli $d\ell e$ are $\le X^{\theta}$ for large X, and the extra choices cost a divisor power. The total probability is at most $o_\eta(1)$ plus a constant times

$$\mathbb{E}_{\mathcal{M}_{s,y}} \sum_{\ell \mid d} \frac{1}{\ell} \ \le \sum_{\ell} \frac{q_\ell}{\ell} + \frac{\sum_\ell c_\ell/\ell}{\sum_\ell c_\ell} = O(s).$$

Here and in this display the model sums run over $2 < \ell \le y$, and $q_\ell \le s\log\ell/(\ell-2)$. Larger squares have dyadic density $O(1/y + X^{-1/2})$ and weight at most $O(L^2)$. Indeed $y \ge \exp(L/\sqrt{u})$ exceeds every power of L. The discarded rows and lists have the bounds already proved.

6 The mark, the least prime and the full factor count

We apply Theorem 1.2 to the independent model (3). The familiar exponential, Poisson and normal limits then transfer jointly to the arithmetic law, including its conditional distributions in mean over D. The choice of the mark and the inclusion of the other primes describe different scales. The designated prime has its own exponential scale. Each other prime has a small inclusion probability, and the accumulated number of inclusions grows like $\log u$.

Theorem 6.1 (Conditional limits from the marked model)**.** *Under* (1)*, consider the quantities*

$$s\log R, \qquad \sum_{\substack{p \mid n+2 \\ p \ne R}} \delta_{s\log p}, \qquad Y = \frac{\Omega(n+2) - \log u}{\sqrt{\log u}}.$$

Test the middle quantity by counts on finitely many disjoint bounded intervals $(a,b] \subset [0,\infty)$. The joint limit consists of an exponential variable of mean one, a Poisson point process of intensity $(1-e^{-v})\,\mathrm{d}v/v$, and a standard normal variable, all independent. For each rectangle with continuity boundaries for these laws, the conditional probability given D converges to the stated probability in mean absolute difference. In particular (5) *and* (6) *hold. Under $\mathbb{Q}_X$, $(\Omega(n), Y)$ converges to $(1 + C_N, Z)$ with independent coordinates, where Z is standard normal.*

Proof. We first work under $\mathcal{M}_{s,y}$. Draw R independently of a pool of Bernoulli variables B_p with probabilities q_p, and erase B_R. The chance that erasure changes the pool is at most

$\max q_p = O(s)$ since $q_p \le s \log p/(p-2)$. The estimate for $\sum c_p$ in Section 5, applied up to $\exp(z/s)$, gives $\Pr(s \log R \le z) \to 1 - e^{-z}$ for each fixed $z \ge 0$. For the other primes,

$$\sum_{2<p\le y} \left| q_p - \frac{1-p^{-s}}{p} \right| \ll s, \qquad \sum_{2<p\le y} q_p^2 \ll s^2 \sum_{p>2} \frac{\log^2 p}{(p-2)^2} \ll s^2.$$

On $a < s \log p \le b$ the sum of the q_p tends to $\int_a^b (1-e^{-v})\,\mathrm{d}v/v$. Indeed $s \sum_{p\le \exp(v/s)} (\log p)/p = v + O(s)$, and the factor $(1-e^{-v})/v$ extends smoothly to zero. Independent Poisson variables of means q_p can replace the individual Bernoulli variables at cost $O(\sum q_p^2)$ in total variation, by coupling each pair [19]. This proves the first two limits and their independence. The event that the least prime exceeds $\exp(z/s)$ requires both the mark to exceed that threshold and the process to have no point up to z. This yields (6).

For the full count put $T = \log u$, $V = T^{1/4}$ and $y = X^{1/V}$. The pool has mean $T - \log V + O(1)$ and variance differing from its mean by $O(s^2)$. To see the mean, the contribution below $\exp(1/s)$ is $O(1)$ by the weighted prime sum just used. Above that point the sum of p^{-1-s} is bounded by partial summation, and Mertens' harmonic prime sum gives $\log(s \log y) + O(1) = T - \log V + O(1)$. Since $\log V = o(\sqrt{T})$, centring by T has the same normal limit. The bounded independent summands satisfy Lindeberg's condition. Removing finitely many bounded v-intervals costs $O(1)$ in expected count. The remaining sum still has the same normal limit and is independent of the counts in the removed intervals and of the mark. Adding the single mark, erasing B_R, and restoring the removed counts change the normalised variable by $o(1)$ in probability.

In the integer experiment, the number of factors $> y$, even with multiplicity, is at most $\log(2X+2)/\log y \le 2V = o(\sqrt{T})$. With $\mathbb{Q}_X$-probability $1 - o(1)$ the neighbour is squarefree by Section 5. Theorem 1.2 transfers the marked counts below y in conditional total variation in mean. The two error events just described have conditional probabilities with the same unconditional means. Bracketing the threshold for Y by its two translates of size $2V/\sqrt{T}$ now proves the conditional rectangle statement. A finite grid of Gaussian quantiles and monotonicity, including both tails, give the supremum in (5). Finally $\Omega(n) = 1 + \Omega(D)$ identically. Its tight marginal limit under $\mathbb{Q}_X$ follows from cofactor transfer and Proposition 2.2. Multiplying a conditional probability by any indicator depending on D does not increase its mean error. This proves the last assertion. □

Corollary 6.2 (Frequency windows)**.** *Let $K = \Omega(n+2)$, $T = \log u$ and $h = h(X) > 0$. If $h\sqrt{T} \to \infty$, then*

$$\sup_{|B_X|\le 1} \mathbb{E}_{\mathbb{Q}_X} \sup_{\alpha\in\mathbb{R}} \frac{1}{2h} \int_{-h}^{h} \left| \mathbb{E}_{\mathbb{Q}_X}(B_X e^{\mathrm{i}(\alpha+v)K} \mid D) \right|^2 \mathrm{d}v \longrightarrow 0.$$

The multiplier B_X may depend on the arithmetic sample but is fixed across the integration window. The scale is sharp. For $B_X = 1$, $\alpha = 0$ and $h\sqrt{T} \to c \in [0,\infty)$ the same integral, before averaging over D, converges in mean absolute difference to $\int_0^1 e^{-c^2 r^2}\,\mathrm{d}r$.

Proof. Take two copies of (B_X, K) independently conditional on D. The modulus of the integrated kernel is at most $\min(1, 1/(h|K-K'|))$, with value one on the diagonal. Write $\epsilon(D)$ for the conditional Kolmogorov distance in (5). For each $d > 0$ the conditional probability of $|K-K'| \le d$ is at most $2d/\sqrt{2\pi T} + 2\epsilon(D)$ by bounding the Gaussian mass of the corresponding interval, using its left limit at the lower endpoint. The complementary kernel costs at most $(hd)^{-1}$. Choose $d = (\sqrt{T}/h)^{1/2}$ and average. More explicitly, with an absolute constant C the left-hand side in the assertion is at most

$$\frac{C}{(h\sqrt{T})^{1/2}} + 2\mathbb{E}_{\mathbb{Q}_X}\epsilon(D).$$

This bound is uniform in both the centre and the arithmetic multiplier.

For the converse put $v = hr$ with $-1 \le r \le 1$. The phase from centring K disappears on taking the modulus. For fixed r, conditional convergence in (5) gives $\mathbb{E}_{\mathbb{Q}_X}(e^{\mathrm{i}rh\sqrt{T}Y} \mid D) \to e^{-r^2c^2/2}$ in mean absolute difference. One may first truncate Y to a compact interval, where the frequencies converge uniformly and integration by parts bounds the error by a constant times $\epsilon(D)$, then use Gaussian tightness for the tails. Squaring the modulus and integrating in r gives the formula by dominated convergence. Hence every bounded subsequential limit of $h\sqrt{T}$ obstructs uniform vanishing. □

The window estimate still allows a large coefficient at one isolated frequency. In particular the single value π, which tests $(-1)^{\Omega(n+2)}$, is an additional arithmetic question. The distinction already occurs for elementary random variables: if $K_m = 2\,\mathrm{Bin}(m, 1/2)$, then $(K_m - m)/\sqrt{m}$ has a standard normal limit, whereas $\mathbb{E}(-1)^{K_m} = 1$ for every m. A normal approximation therefore leaves this isolated frequency undetermined.

A Distribution in progressions with the full count retained

The deletion argument used rough first-site factors to allow an arbitrary test of the cofactor. This appendix gives an independent application of Granville and Shao's theorem [15, Theorem 2.1]: it keeps the full divisor weight and labels only the count. Its conclusion is a uniform specialisation of that theorem and is not needed for the cofactor transfer. Write $\Delta(v; q, a)$ for the sum of v on $X < n \le 2X$, $n \equiv a \bmod q$, less the sum on integers coprime to q in the same interval divided by $\varphi(q)$.

Proposition A.1. *For $A, J > 0$ there exists $B > 0$ such that, uniformly for $s \ge 0$ and all complex functions $|b_X(k)| \le 1$,*

$$\sum_{q \le \sqrt{X}/L^B} \tau(q)^J \max_{(a,q)=1} |\Delta(b_X(\Omega)W_s; q, a)| \ll_{A,J} XL^{-A}.$$

Proof. Begin with $f(n) = z^{\Omega(n)}H_s(n)$, $|z| = 1$. Its local Dirichlet factor at p is $(1 - zp^{-w-s})/(1 - zp^{-w})$. Its logarithmic derivative coefficients at p^j are $z^j(1 - p^{-js})\log p$, bounded in modulus by $\Lambda(p^j)$. Thus f belongs to the class $\mathcal{C}$ of Granville and Shao in *When does the Bombieri–Vinogradov theorem hold for a given multiplicative function?* [15, Theorem 2.1].

For clarity, fix a desired exponent $A_0 > 2$ in that theorem, choose $\gamma = 2A_0 + 7$, and put $y_0 = x/(\log x)^\gamma$, $Q = \sqrt{x}/(\log x)^B$. The prime restriction of f is $z(1 - p^{-s})$, a multiplier of bounded variation uniformly in s. Maximal Bombieri–Vinogradov and partial summation give the required prime discrepancy, summed up to Q, of at most $O(Y/((\log x)^{A_0}\log(x/y_0)))$ for $y_0 \le Y \le x$. Indeed choose $B > A_0 + 5$ with an additional margin exceeding $\gamma/2$ plus the level loss of the prime theorem at the desired accuracy. This choice also accommodates either endpoint of our dyad.

The remaining hypothesis concerns individual moduli, with error $O(Y(\log x)^{-(A_0+2B)})$ for $\sqrt{x} \le Y \le x$. Here are details of the uniformity in s, z. For a nonprincipal character $\chi \bmod q$, $q \le (\log x)^D$, the twisted series is

$$L(w, \chi)^z L(w + s, \chi)^{-z} E(w).$$

The residual logarithm is bounded uniformly for $\Re w \ge 3/4$, since only terms of degree at least two remain. Take the analytic logarithms from $\Re w > 1$. Up to height $2\exp(\sqrt{\log x})$ there is a zero-free strip of width $c_D/\sqrt{\log x}$ at 1 for these characters. Siegel's bound $1 - \beta \gg_\epsilon q^{-\epsilon}$ handles a possible real exceptional zero on taking $\epsilon < 1/(2D)$. In a slightly smaller strip $\log L(w, \chi) = O_D(\log\log x)$. Indeed finite truncation gives $|L| \ll_D (\log x)^{O(1)}$ there (partial

sums of a nonprincipal character are $O(q)$). On the line $1 + c_D/\sqrt{\log x}$ the Euler product bounds the full logarithm by $O(\log\log x)$. Borel–Carathéodory in zero-free discs centred on this line gives the asserted bound after shrinking the strip. This applies to imprimitive characters as well, whose missing Euler factors are zero-free there. It also applies to $w + s$, with the Euler product available further to the right. These classical zero-free and Perron estimates may be found in *Analytic number theory* [18, Chs. 5 and 18].

Perron's formula for bounded coefficients, truncated at height $\exp(\sqrt{\log x})$ and moved to the smaller strip, now saves every fixed power of $\log x$ on $\sqrt{x} \le Y \le x$. One may take the endpoint half-integral and bound the truncation error by $O(1 + Y(\log x)^2 e^{-\sqrt{\log x}})$. Character orthogonality gives the desired individual discrepancy. For $q > (\log x)^D$ the elementary bound $O(Y/q + 1 + Y/\varphi(q))$ suffices on choosing D after A_0, B. Granville and Shao therefore give the summed discrepancy for $z^\Omega H_s$ with bound $O(x/(\log x)^{A_0 - 1})$. The choices are independent of s, z; the constants may be ineffective because the argument uses Siegel's bound.

Finally, differentiating the finite product and using that all undifferentiated factors lie in $[0, 1]$ gives, for $s \ge 0$,

$$\begin{aligned} |H_s''(n)| &\le \sum_{p \mid n} (\log p)^2 + 2 \sum_{\substack{p<q \\ p,q \mid n}} (\log p)(\log q) \\ &= \left(\sum_{p \mid n} \log p \right)^2 \le (\log n)^2. \end{aligned}$$

A forward difference of increment L^{-M} therefore recovers W_s to pointwise error $O(L^{2-M})$, uniformly in s. Choose M for the required error after summing the elementary progression bounds, then choose the accuracy for H_s larger by M. For an integer $k_0 > \log(2X)/\log 2$ of order L, Fourier expansion on $\mathbb{Z}/k_0\mathbb{Z}$ represents $b_X(\Omega)$ by at most k_0 unimodular powers with coefficients bounded by one. Absorb this logarithm in the initial accuracy. To insert τ^J, apply Cauchy–Schwarz as in Section 3, bounding the sum weighted by τ^{2J} trivially by $XL^{O_J(1)}$ since $|W_s(n)| \le \log n$. This proves the proposition. □

B Marked factors in binary and ternary additive equations

Suppose now that every parameter t_i is fixed and that all sites are sampled together subject to one linear equation. For binary targets one can obtain the marked law at both sites outside a small exceptional set of targets. A third summand removes that exceptional set. We include the argument because the population that is averaged is essential to its interpretation.

Take $k = 2$ or 3, signs $e_i = \pm 1$, fixed compact intervals $I_i \subset (0, \infty)$ of positive length, and $t_i \ge 0$. Give solutions $n_i \in XI_i$ of $\sum e_i n_i = h$ weight $\prod_i W_{t_i/L}(n_i)$. Mark one prime independently at each site conditional on the integers, using the positive summands of its weight, and order its other distinct primes uniformly. Set

$$J = \mathbf{1}_{e_1 I_1} * \cdots * \mathbf{1}_{e_k I_k}, \qquad \mathfrak{S}_k(h) = \sum_{q \ge 1} \frac{\mu(q)^k c_q(h)}{\varphi(q)^k}, \quad c_q(h) = \sideset{}{^*}\sum_{a \bmod q} e^{2\pi i a h/q},$$

where the star restricts to units and μ is the Möbius function. Call a target admissible if $h \equiv k \bmod 2$, $J(h/X) \ge c_0 > 0$, and $h \ne 0$ for $k = 2$.

Theorem B.1. *At the targets specified below the total weight, divided by $X^{k-1} J(h/X) \mathfrak{S}_k(h)$, tends to one. The normalised law of the marked lists of logarithmic sizes $\log p/L$ tends weakly to $\bigotimes_i \mathcal{P}_{t_i}$, and the vector of counts $\Omega(n_i)$ tends in total variation to its count marginal.*

1. *For $k = 2$ and $0 < |h| \le CX$, the number of exceptional admissible targets is $O_{A,\delta}(X/L^A)$ for every $A, \delta > 0$.*

2. *For $k = 2$, fix $a \in \mathbb{Z} \setminus \{0\}$, $d \geq 1$ integral, and $B \geq 0$. Take $h = b_X + ap^d$ with $b_X \in \mathbb{Z}$ and $|b_X| \leq C_b X$ for a fixed $C_b > 0$, $Y < p \leq 2Y$ prime and $X/L^B \ll Y^d \ll X$. There are $O_{A,\delta}(Y/(\log Y)^{A+1})$ exceptional prime indices with admissible target.*

3. *For $k = 3$ the assertions are uniform in all admissible h.*

In the first two statements an exception means an error exceeding δ for a fixed bounded continuous size test, the relative total mass, or the total variation distance for the count vector. Constants may depend on the fixed data and on the size test. In the second statement they are uniform in b_X subject to the displayed bound.

For example, take $e_1 = e_2 = e_3 = 1$, $X = h$ odd, with $1/3$ in the interior of each I_i. The theorem gives a product law at each such h. For the signs $+, +, -$, one can even fix $h = 1$, $I_1 = I_2 = [1, 2]$, $I_3 = [2, 4]$, for which $J(1/X) \to 1$. For the difference of two sites with overlapping intervals and $h = 2p^d$, the binary singular series tends to $\mathfrak{S}_2(2)$ as $p \to \infty$. The exceptional-set statement still allows the single shift $h = 2$ to be exceptional. The binary representation of sparse sequences by sums of primes is classical. Perelli's *Goldbach numbers represented by polynomials* [22] and Brüdern, Kawada and Wooley's *Additive representation in thin sequences, II: The binary Goldbach problem* [5] treat polynomial targets; Brüdern and Perelli [6] treat sparse sequences more generally. For a fixed polynomial, sufficiently small exceptional sets of integer indices already give small exceptional sets of prime indices. The statement here tracks critical divisor weights, designated primes and joint factorisation laws, with the stated uniformity for the moving translation b_X. This appendix is an additive complement to the fixed-shift observation results.

Removing small factors uniformly in the target

Fix $0 < \eta < 1/10$. As in Lemma 2.1, use $W_{t_i/L}(n_i) \ll Lf(n_i)$ with $f(p^j) = \min(1, K \log p/L)$ and K a fixed sufficiently large constant. Force $p \leq X^\eta$ to divide the first site and write $n_1 = pm$. Take m and, for $k = 3$, n_2 as free variables. Enlarge their domain to a rectangle by dropping the interval condition on the final site. The short side has length comparable with $U = X/p$ and the volume is $O(X^{k-1}/p)$. Remove all p-factors except the one factor $f(p)$ that has been extracted at the forced site.

Retain each prime $2k < \ell \leq U^{1/(40k)}$, $\ell \neq p$, independently with probability $1 - f(\ell)$ and forbid it to divide the product of the k forms. At an integer point, average survival bounds the desired product of f's from above, including when a prime divides several forms. For two sites the forbidden proportion is $g(\ell) = \nu_\ell/\ell$ with $\nu_\ell = 1$ if $\ell \mid h$ and 2 otherwise. For three sites the forbidden sets are three lines of distinct directions, so $g(\ell) = 3/\ell + O(1/\ell^2)$ uniformly in h, p. On squarefree moduli $d \leq U^{2/5}$ the actual proportion in the residue sets is $g(d) + O(k^{\omega(d)}/U)$, by multiplicativity and counting points in each class with relative error $O(d/U)$. Here $dg(d) \leq k^{\omega(d)}$ for the retained primes. In the upper Selberg sieve with cutoff $U^{1/5}$ the summed remainder is $O(U^{-1/2})$, just as in Lemma 2.1. The expected log of the divisor in the full Euler denominator is at most $k \log(U^{1/(40k)}) + O(1)$, so the truncated denominator retains a fixed fraction by Markov's inequality. On averaging the main product and using $\sum_{\ell \ll X} f(\ell)/\ell = O(1)$ one obtains $O(L^{-3})$ for three sites and $O(L^{-2}\mathfrak{S}_2(h))$ for two. For the latter, the excess at $\ell \mid h$ is bounded by the product $\prod_{\ell \mid h,\ \ell > 2}(\ell - 1)/(\ell - 2)$.

Restore $L^k f(p)$, use the bound $\sum_{p \leq X^\eta} f(p)/p \ll \eta + 1/L$, and repeat for each site. The removed weight, divided by X^{k-1}, is $O((\eta + 1/L)\mathfrak{S}_2(h))$ for two sites and $O(\eta + 1/L)$ for three. A square of a larger prime at a prescribed site costs $O(X^{k-1}L^k(X^{-\eta} + X^{-1/2}))$, by taking that site free and using $W_s(n) \leq \log n$. This also proves tightness of the factor counts after the cutoff is removed.

Counting logarithmic cells

Fix now the number r of distinct factors at a site, all $> X^{\eta}$. A cell requires $x_i = \log p_i/L \in J_i$ for the designated prime p_0 and the ordered remaining primes. Use intervals with boundaries of limiting mass zero, in particular avoiding the designated atom at one. Define $A(n)$ to be the sum over $p_0 \cdots p_{r-1} = n$ in the cell, $n \in XI$, of

$$\frac{(\log p_0)e^{-tx_0}}{(r-1)!}\prod_{i>0}(1-e^{-tx_i}).$$

Dropping distinctness costs $O_\eta(X^{1-\eta}L + X^{1/2}L)$ in the one-site sum, because the number of tuples per integer is bounded after truncation. Put $T(\alpha) = \sum_n A(n)e(\alpha n)$, $e(w) = \exp(2\pi \mathrm{i} w)$. For $v \in I$ define

$$g_X(v) = \frac{1}{(r-1)!}\int \mathbf{1}_{J_0}(x_0)e^{-tx_0}\prod_{i>0}\mathbf{1}_{J_i}(x_i)(1-e^{-tx_i})\,\frac{\mathrm{d}x_i}{x_i}, \qquad x_0 = 1 + \frac{\log v}{L} - \sum_{i>0} x_i,$$

and take $g_X = 0$ off I. Its supremum and variation are bounded uniformly in X by integration of translates of bounded-variation functions. On I it tends uniformly to the cell mass in (8). For $r = 1$ this uses the choice of endpoints. For $r > 1$ it follows by continuity in translation of the bounded integrable densities on the truncated intervals.

For $B_1 = L^{K_1}$ let the major arcs be $|\alpha - a/q| \le B_1/X$, $q \le B_1$, $(a,q) = 1$, on the unit circle. For each desired A_0, K_1 can be chosen so that

$$\sup_{\text{minor}} |T(\alpha)| \ll XL^{-A_0},$$
$$T(a/q+\beta) = \frac{\mu(q)}{\varphi(q)}X\widehat{g}_X(X\beta) + O_{A_1,K_1}(XL^{-A_1}) \tag{15}$$

for any A_1, where $\widehat{g}_X(\xi) = \int g_X(v)e(\xi v)\,\mathrm{d}v$. On the major arcs each prime exceeds q. Split into unit classes for all variables. At fixed values of the other variables the product cutoff gives an interval in the remaining one. Siegel–Walfisz and partial summation replace each prime sum successively by integration against $\mathrm{d}p/(\varphi(q)\log p)$. The phase has variation $O(1+B_1)$ on each such interval. There are at most q^r residue tuples and $O(L^r)$ dyadic boxes of product scale $O(X)$, so arbitrary saving survives. The product of uniform units is uniform, giving the displayed Ramanujan coefficient. Substituting $p_0 = Xv/\prod_{i>0} p_i$ gives the asserted g_X.

For $r = 1$ the classical minor-arc bound for a prime sum, with partial summation, gives the first estimate in (15) [18, Ch. 13]. For $r \ge 2$ group the factors into two nonempty sets of product scales $M, N \gg X^{\eta}$, $MN \asymp X$. Subdivide both ranges into intervals of relative width L^{-K_0}. Boundary products from the constraint $n \in XI$ cost $O_\eta(XL^{1-K_0} + L)$ since each coefficient is $O_\eta(L)$. The bin coefficients are separately bounded by $O_\eta(L)$ from boundedly many factorizations. For $|\alpha - a/q| \le q^{-2}$ the resulting bilinear sum on each box, divided by MN, is bounded by

$$L^{O(1)}(q^{-1} + M^{-1} + N^{-1} + q/(MN))^{1/2}.$$

Indeed use Cauchy–Schwarz in the first variable, expand, and bound the geometric progression at each difference j in the second variable by $\min(M, \|j\alpha\|^{-1})$. Summing over $1 \le j \ll N$ costs $O((N/q+1)(M + q\log(2q)))$, by spacing in blocks of length at most $q/2$. Off the major arcs Dirichlet approximation gives $B_1 < q \le X/B_1$ up to harmless integer rounding. Choose K_0 first for the boundaries, and K_1 subsequently to absorb all boxes. This proves (15).

From cells to targets

Write $G_X = g_{1,X}^{e_1} * \cdots * g_{k,X}^{e_k}$ with $g^e(v) = g(ev)$. It tends uniformly to J times the product of the one-site cell masses. Retain G_X until the logarithmic error estimates are completed:

its convergence to the limiting cell profile supplies only $o(1)$, which would not suffice for the exceptional-set bounds $O_A(X/L^A)$. With $k=2$ let $E(h)$ be the binary cell mass divided by X less $G_X(h/X)\mathfrak{S}_2(h)$, and set $E(0)=0$. For every $A>0$,

$$X^{-1}\sum_{|h|\ll X}|E(h)|^2 \ll_A L^{-A}. \tag{16}$$

Indeed Parseval on the minor-arc product uses one supremum in (15) and the second-moment bound $\sum_n |A(n)|^2 \ll XL^2$. On the major arcs use the second line of (15) with arbitrarily large A_1. The singular integral extends to all real frequencies at cost $O(B_1^{-1})$ by bounded variation of the two profiles, since each transform is $O((1+|\xi|)^{-1})$. The absolute sum of the rational coefficients up to B_1 is $O(\log(2B_1))$ by $|c_q(h)|\le\varphi(q)$. Beyond B_1 their absolute sum has mean over nonzero h of size $O(\sum_{q>B_1}\mu(q)^2\tau(q)/\varphi(q)^2)\ll B_1^{-1/2}$, using $|c_q(h)|\le\sum_{a|(q,h)}a$ and counting nonzero multiples. The tail is uniformly $O(L)$ because its full absolute Euler product is $O(\prod_{\ell|h}(1+1/(\ell-1)))=O(L)$ on this range. This yields (16) after increasing K_1.

For $k=3$, bound one minor-arc sum by its supremum and the other two by Parseval and Cauchy–Schwarz. The error after division by X^2 is $O_A(L^{-A})$ for each target. The series tail is uniformly $O(B_1^{-1/2})$ by $|c_q(h)|\le\varphi(q)$, and the singular integral extends as above. At odd h the Euler factor at two is two. At an odd prime it is $1-1/(p-1)^2$ if $p\mid h$, and $1+1/(p-1)^3$ otherwise. This gives uniform upper and positive lower bounds for $\mathfrak{S}_3(h)$.

To pass from (16) to sparse targets, embed all the signed supports without wraparound in a cyclic group of prime order $H\asymp X$. Use the normalised discrete Fourier transform. We claim $\|\widehat{E}\|_1\ll L^{O(1)}$. For the physical convolution this follows from Cauchy–Schwarz and Parseval on the two transforms with $\sum|A_i|^2\ll XL^2$. For the model use, at nonzero h,

$$\mathfrak{S}_2(h)=\mathfrak{S}_2(2)\mathbf{1}_{2|h}\sum_{\substack{a|h\\ a\text{ odd}}}\mu(a)^2\prod_{p|a}(p-2)^{-1}.$$

The positive coefficients sum to $O(L)$ for $a\ll X$ by the Euler product. In a signed interval of length $O(X)$, multiples of $2a$ form a progression with Fourier ℓ^1 norm $O(\log H)$ by the geometric sum. Partial summation preserves that bound with bounded-variation multiplier G_X. Correcting the value at zero costs at most $O(L^2)$ in the Fourier norm. This proves the claim. For $D_0=|E|^2$, convolution and (16) now give

$$\|\widehat{D}_0\|_1\ll L^{O(1)},\qquad \|\widehat{D}_0\|_\infty\ll_A L^{-A}\quad\text{for every } A.$$

The probability measure on b_X+ap^d, $Y<p\le 2Y$, with masses proportional to $\log p$, has Fourier multiplier given by the normalised sum of $e((b_X+ap^d)\xi/H)\log p$. By the prime polynomial minor-arc bound in Chapman and Chow's *Arithmetic Ramsey theory over the primes* [7, Lemma 4.2], applied at both endpoints with constant amplitude, this is $O_{A'}(L^{-A'})$ unless some $q\le L^{K_2}$ satisfies $Y^d\|qa\xi/H\|\le L^{K_2}$, with K_2 depending on A',d. Multiplication by qa permutes residues modulo H for large X. Hence there are $O(L^{K_2}(1+HL^{K_2}/Y^d))=L^{O(1)}$ exceptional frequencies. In the Fourier average of D_0 use its supremum bound there and its ℓ^1 bound elsewhere. This gives mean square $O_A(L^{-A})$ also over the prime indices, since the dyadic logarithmic weights are comparable. Chebyshev gives both binary exceptional-set estimates.

Finally choose η for the prescribed tolerance. The preceding roughness bounds are relative to the total main term since $J(h/X)\ge c_0$ and $\mathfrak{S}_2(h)\ge\mathfrak{S}_2(2)>0$ for nonzero even h. On the truncated populations there are finitely many factor counts. Finite partitions by continuity cells give their joint masses, and approximate every fixed continuous size test on each compact truncated stratum. Apply the cell estimates to the finite partition and use the one-site tail bound in Proposition 2.2 for the limiting model. This proves Theorem B.1, including total

variation for the count vector. For instance, testing each site by $(-1)^{\Omega}$ gives the limiting product $\prod_i (t_i - 1)e^{-t_i}$. On a uniform permutation of m letters the mean of $(-1)^{C_m}$ is zero for $m \geq 2$ by pairing opposite permutation signs, and the cases $m = 0, 1$ have means $1, -1$.

Matching the local data of the shift two

There is a sequence of growing shifts that has both the singular series and every fixed congruence of the shift two. More precisely, take $n, n + h_X \in (X, 2X]$. One can choose $h_X \to \infty$, $h_X/X \to 0$ and $h_X \equiv 2$ modulo every fixed modulus for all sufficiently large X, such that locally uniformly in $z \in \mathbb{C}$,

$$\frac{1}{X} \sum_{n, n+h_X \in I} W_{z/L}(n) W_{z/L}(n + h_X) \longrightarrow S,$$

$$\frac{1}{X} \sum_{n, n+h_X \in I} \lambda(n)\lambda(n + h_X) W_{z/L}(n) W_{z/L}(n + h_X) \longrightarrow S(z-1)^2 e^{-2z},$$

where $\lambda(n) = (-1)^{\Omega(n)}$ and $I = (X, 2X]$. Complex weights here mean analytic extensions of finite sums.

To construct the shifts, let $Y^d = X/L$ for a fixed $d \geq 1$, put $M_X = \operatorname{lcm}(1, \ldots, \lfloor \log\log X \rfloor)$, and take $h = 2p^d$, $Y < p \leq 2Y$, $p \equiv 1 \bmod M_X$. The prime number theorem gives $M_X = L^{1+o(1)}$. Siegel–Walfisz supplies $(1 + o(1))Y/(\varphi(M_X) \log Y)$ candidates. At any fixed tolerance, the exceptional set in Theorem B.1 for finitely many real $t_i = z \geq 0$ is a vanishing fraction of this set. All candidates satisfy $J(h/X) \to 1$ and $\mathfrak{S}_2(h) = S(p-1)/(p-2) \to S$. Diagonal selection through rational $z \geq 0$ and decreasing tolerances gives a single h_X with convergence for both observables on this dense set.

For $|z| \leq R$ one has $|W_{z/L}(n)| \ll_R L f_R(n)$, where $f_R(p^j) = \min(1, K_R \log p/L)$ with $K_R > 2R + 2$. To see this, use $|1 - e^{-zx}| \leq e^{Rx} \min(Rx, 2)$ in each unmarked factor. The product of the exponential costs is bounded because $\sum_{p|n} \log p/L \leq 1 + O(1/L)$. There are $O_R(1)$ factors with $K_R x \geq 1$. For the smaller ones, division by $f_R(p)$ leaves a factor at most R/K_R when unmarked, and at most L/K_R for the logarithm of a mark. Summing over all possible marks gives the bound as in Lemma 2.1. The binary upper sieve above, now without prescribing p, bounds the average of $L^2 f_R(n) f_R(n + h_X)$ by $O_R(\mathfrak{S}_2(h_X))$. The two entire functions of z are therefore locally bounded. Normal-family compactness and uniqueness of analytic continuation extend the limits to every compact set. The construction retains the freedom to choose a prime index at every scale. Control at the fixed shift two remains open for these binary product laws.

Acknowledgement

OpenAI's ChatGPT assisted with the mathematical exploration and preparation of this article. Its Codex environment was used to develop the row-by-row progression estimate and the stopping-rule argument, check bibliographic attributions, and implement finite checks of the sampling identities. OpenAI's GPT-6 Astra Pro also assisted with mathematical, bibliographic and editorial revisions. The author retains responsibility for the manuscript. The source archive `Visible_Prime_Factors_sources.zip` includes the exact finite checks and their reference output. From its root, `python3 checks.py` replays them, including the comparison of the exact and censored mark laws.

References

[1] R. Arratia, *On the amount of dependence in the prime factorization of a uniform random integer*, in *Contemporary Combinatorics*, Bolyai Soc. Math. Stud. **10**, János Bolyai Math. Soc., Budapest (2002), 29–91. Archived text: arXiv:1305.0941v1.

[2] R. Arratia and F. Kochman, *A simple direct proof of Billingsley's theorem*, arXiv:1401.1553 (2014). 10.48550/arXiv.1401.1553.

[3] P. Billingsley, *On the distribution of large prime divisors*, Period. Math. Hungar. **2** (1972), 283–289. 10.1007/BF02018667.

[4] E. Bombieri, *On twin almost primes*, Acta Arith. **28** (1975), 177–193. 10.4064/aa-28-2-177-193. *Corrigendum to my paper "On twin almost primes" and an addendum on Selberg's sieve*, ibid. **28** (1976), 457–461. 10.4064/aa-28-4-457-461.

[5] J. Brüdern, K. Kawada and T. D. Wooley, *Additive representation in thin sequences, II: The binary Goldbach problem*, Mathematika **47** (2000), 117–125. 10.1112/S002557930001576X.

[6] J. Brüdern and A. Perelli, *Goldbach numbers in sparse sequences*, Ann. Inst. Fourier **48** (1998), 353–378. 10.5802/aif.1621.

[7] J. Chapman and S. Chow, *Arithmetic Ramsey theory over the primes*, Proc. Roy. Soc. Edinburgh Sect. A **156** (2026), 966–1012. 10.1017/prm.2024.96.

[8] D. Elboim and O. Gorodetsky, *Multiplicative arithmetic functions and the generalized Ewens measure*, Israel J. Math. **262** (2024), 143–189. 10.1007/s11856-024-2609-x.

[9] W. J. Ewens, *The sampling theory of selectively neutral alleles*, Theoret. Population Biol. **3** (1972), 87–112. 10.1016/0040-5809(72)90035-4.

[10] K. Ford, *Poisson approximation of prime divisors of shifted primes*, Int. Math. Res. Not. IMRN (2025), no. 7, rnaf079. 10.1093/imrn/rnaf079. Corrected text: arXiv:2408.03803v4, 19 January 2026.

[11] J. Friedlander and H. Iwaniec, *On Bombieri's asymptotic sieve*, Ann. Scuola Norm. Sup. Pisa Cl. Sci. (4) **5** (1978), 719–756. Numdam: ASNSP_1978_4_5_4_719_0.

[12] J. Friedlander and H. Iwaniec, *Opera de cribro*, Amer. Math. Soc. Colloq. Publ. **57** (2010). 10.1090/coll/057.

[13] O. Gorodetsky, *A Kubilius model for sieve-theoretic sequences*, Anal. Math. **52** (2026), 999–1022. 10.1007/s10476-026-00178-w. arXiv:2608.14190v1.

[14] A. Granville, *Prime divisors are Poisson distributed*, Int. J. Number Theory **3** (2007), 1–18. 10.1142/S1793042107000778. Erratum, ibid., 649–651. 10.1142/S1793042107001073.

[15] A. Granville and X. Shao, *When does the Bombieri–Vinogradov theorem hold for a given multiplicative function?*, Forum Math. Sigma **6** (2018), e15. 10.1017/fms.2018.14.

[16] B. Green and T. Tao, *Restriction theory of the Selberg sieve, with applications*, J. Théor. Nombres Bordeaux **18** (2006), 147–182. 10.5802/jtnb.538.

[17] T. Haddad and D. Koukoulopoulos, *On Arratia's coupling and the Dirichlet law for the factors of a random integer*, J. Éc. polytech. Math. **12** (2025), 1565–1604. 10.5802/jep.317.

[18] H. Iwaniec and E. Kowalski, *Analytic number theory*, Amer. Math. Soc. Colloq. Publ. **53** (2004). 10.1090/coll/053.

[19] L. Le Cam, *An approximation theorem for the Poisson binomial distribution*, Pacific J. Math. **10** (1960), 1181–1197. 10.2140/pjm.1960.10.1181.

[20] Y. Motohashi, *An induction principle for the generalization of Bombieri's prime number theorem*, Proc. Japan Acad. **52** (1976), 273–275. 10.3792/pja/1195518296.

[21] M. Nair and G. Tenenbaum, *Short sums of certain arithmetic functions*, Acta Math. **180** (1998), 119–144. 10.1007/BF02392880.

[22] A. Perelli, *Goldbach numbers represented by polynomials*, Rev. Mat. Iberoam. **12** (1996), 477–490. 10.4171/RMI/205.

[23] D. H. J. Polymath, *Variants of the Selberg sieve, and bounded intervals containing many primes*, Res. Math. Sci. **1** (2014), article 12. 10.1186/s40687-014-0012-7. arXiv:1407.4897v4, 22 December 2014.

[24] B. Pouly, *Critical divisor weights and the localisation of parity defects*, Zenodo (2026). 10.5281/zenodo.24578863. Theorem 1.1 and (2).

[25] B. Pouly, *Observations and Shift Decoding in Divisor Incidences*, Zenodo (2026). 10.5281/zenodo.24578881. Theorem 1.2.

[26] B. Pouly, *Reconstruction and sign prediction from sampled prime factors*, Zenodo (2026). 10.5281/zenodo.24443266.

[27] W. Sun, *Tilting Billingsley's model toward a giant prime: two phase transitions*, preprint, 10 September 2026. arXiv:2609.11735v1.

[28] S. Tavaré, *The magical Ewens sampling formula*, Bull. London Math. Soc. **53** (2021), 1563–1582. 10.1112/blms.12537.