Prime noise, polynomial observations and trace-preserving completion of multiplicative functions

Brice Pouly

Research and writing with AI assistance

hexagon:2610.00351v2

Abstract

Given a completely multiplicative sign function and an eventually positive integer polynomial irreducible over $\mathbb{Q}$ of degree at least two, we construct a random completion preserving every value on the polynomial trace and any prescribed finite prefix. Almost surely every fixed admissible affine correlation of the completion vanishes at every positive time. Independent clocks on primes that cannot divide the trace provide the construction. When all primes carry clocks, the irreducible factors of the observed rational functions determine the local transition law. We then study what can be detected after the signs are drawn. On products of two primes, separately chosen row and column weights still give cancellation under von Mangoldt sampling. One shared sign per band can also preserve every value on n^2+1; the width $\log 2$ is the sharp summability threshold for prescribed tables varying across dyadic scales. On labelled divisor cells one cofactor bit transmits a fraction $e^{-\theta/2}$ of the initial bias at the critical cutoff. Twists of finite quadratic prime distance preserve cancellation on the family of fixed odd dilations. For ordinary intervals with independent flips on a moving band, bounded periodic tests retain the degree-d multiplier e^{-2dt} uniformly up to every period bound $o(X)$. Two fluctuation experiments explain when a signed mean can be recovered from typical draws.

1 Preserving a trace while regularising other observations

Normality of a sign sequence concerns words of fixed length. A trace such as $n \mapsto f(n^2+1)$ samples a multiplicative function in a different way. We ask whether the entire trace of a prescribed source can survive inside a function with vanishing affine correlations.

For a concrete example, a prime $p \equiv 3 \pmod 4$ never divides n^2+1. Flipping signs only at these primes, beyond a prescribed prefix, therefore preserves every value of this trace. Such a prime does act on a non-constant affine form whose leading coefficient it does not divide: the form has one root modulo p. Decreasing clock rates can thus leave the polynomial trace untouched while supplying enough harmonic mass to regularise affine observations. With the Liouville source $f = \lambda$, all values $\lambda(n^2+1)$ survive exactly. The affine conclusion concerns the constructed function. The trace's mean remains unknown.

There is a second way for an observation to survive. Someone who sees the prime signs can try to undo the randomisation. On products of two primes, two separately chosen weights still leave cancellation, while a single joint weight can reproduce every sign. On a labelled divisor progression, even one carefully chosen bit of a cofactor recovers a positive fraction of the original bias. On an ordinary interval a bounded residue weight can follow the draw throughout a sharp range of periods. In each case the order of observation determines how much of the source remains measurable.

1.1 The function classes and the completion theorem

Write $\mathcal{M}_\pm$ for the completely multiplicative functions $f : \mathbb{N}_{\geq 1} \to \{-1, 1\}$. Each extends to $\mathbb{Q}_{>0}$ by $f(a/b) = f(a)f(b)$ and takes the value one on positive rational squares.

An admissible affine family is a fixed non-empty family

$$\ell_i(n) = A_i n + B_i, \qquad A_i \in \mathbb{N}_{\geq 1}, \quad B_i \in \mathbb{Z}, \qquad A_i B_j - A_j B_i \neq 0 \quad (i \neq j). \tag{1.1}$$

Let $\mathcal{L}_\pm$ consist of the $g \in \mathcal{M}_\pm$ such that, for every family (1.1),

$$\frac{1}{X} \sum_{n \leq X} \prod_{i=1}^{k} g(\ell_i(n)) \longrightarrow 0. \tag{1.2}$$

A fixed initial segment is omitted when some arguments are not positive. This does not change the limit. Expanding word indicators over all subfamilies of consecutive shifts gives frequency 2^{-m} for each word of length m. Thus every member of $\mathcal{L}_\pm$ is normal.

Theorem 1.1 (Completion of an irreducible trace)**.** *Let $P \in \mathbb{Z}[T]$ be irreducible over $\mathbb{Q}$, of degree at least two, and positive at all sufficiently large positive integers. For every $f \in \mathcal{M}_\pm$ and every integer $M_0 \geq 1$, there exists $g \in \mathcal{L}_\pm$ such that*

$$g(P(n)) = f(P(n)) \qquad (n \in \mathbb{Z},\ P(n) > 0), \tag{1.3}$$

$$g(n) = f(n) \qquad (1 \leq n \leq M_0), \tag{1.4}$$

$$\frac{\#\{p \leq x : g(p) \neq f(p)\}}{\pi(x)} \longrightarrow 0. \tag{1.5}$$

More precisely, there is a process obtained from independent prime Poisson clocks, starting from f, for which these conclusions and membership in $\mathcal{L}_\pm$ hold almost surely at all real times $t > 0$ simultaneously. The preservation assertions (1.3)–(1.4) *hold at every time, including $t = 0$.*

The source is fixed before the clocks. No initial balance, normality or convergence of its correlations is assumed. The theorem preserves a trace *already compatible with a multiplicative function*; it does not prescribe arbitrary signs on the values of P.

With $\mathrm{Tr}_P(f) = (f(P(n)))_{n \in \mathbb{Z},\, P(n) > 0}$, we obtain

$$\mathrm{Tr}_P(\mathcal{L}_\pm) = \mathrm{Tr}_P(\mathcal{M}_\pm). \tag{1.6}$$

Here "completion" refers to the protected data, not to missing values of f: the source is already defined on every integer, and we realise its protected data inside the smaller class $\mathcal{L}_\pm$. Equality (1.6) does not say that every binary sequence is a compatible trace. It fails for a positive affine polynomial: the trace of the constant source $f \equiv 1$ would contradict the one-form instance of (1.2).

For probability measures on a common measurable space we use the convention

$$d_{\mathrm{TV}}(\mu, \nu) = \sup_{B \text{ measurable}} |\mu(B) - \nu(B)|.$$

On a finite state space this equals $\frac{1}{2} \sum_x |\mu\{x\} - \nu\{x\}|$. The next statement uses this distance for the joint law of the affine signs and the protected marks.

Corollary 1.2 (Preserved marks)**.** *In the construction of Theorem 1.1, fix an admissible affine family $(\ell_i)_{i \leq k}$ and a finite-alphabet sequence $z(n)$, fixed before the clocks and formed from source values at protected arguments. For example,*

$$z(n) = \big(f(P(n + b_1)), \ldots, f(P(n + b_r))\big)$$

for fixed integer shifts. Choose an integer $n_0 \geq 0$, fixed before the clocks and large enough that all arguments under consideration are positive, and set

$$I_X = \{n_0 + 1, \ldots, n_0 + X\}, \qquad \nu_X = \frac{1}{X} \sum_{n \in I_X} \delta_{z(n)}.$$

Almost surely, for every $0 < u < T < \infty$,

$$\sup_{u \leq t \leq T} d_{\mathrm{TV}} \left(\frac{1}{X} \sum_{n \in I_X} \delta_{((F_t^f(\ell_i(n)))_{i \leq k},\, z(n))}, \mathrm{Unif}(\{-1,1\}^k) \otimes \nu_X \right) \longrightarrow 0. \tag{1.7}$$

The same assertion holds separately for $I_X = [X, 2X) \cap \mathbb{Z}$, for all sufficiently large X. No convergence of ν_X is required.

Here F_t^f denotes the constructed function at time t. For one affine form, the joint mass of either sign and a prescribed mark state z is $\frac{1}{2}\nu_X\{z\} + o(1)$, uniformly on the stated time intervals. This is an absolute error; it does not give a relative conditional estimate when $\nu_X\{z\}$ tends to zero.

The probability-one event may depend on the fixed marks; a countable collection of mark sequences can be treated on one event. These are empirical laws as the origin ranges over the specified intervals, not independence of all integer positions. In particular, no uniformity over all translates of an interval, or over arbitrary moving origins, is asserted.

1.2 Two noise regimes and two scales

For the completion and polynomial-channel results of Sections 2–8, the model is

$$W_t(r) = (-1)^{\sum_p v_p(r) Q_p(t)}, \qquad F_t^f(r) = f(r) W_t(r), \qquad r \in \mathbb{Q}_{>0}, \tag{1.8}$$

where the Q_p are independent Poisson processes of rates $a_p \in [0,1]$, starting at zero. Each rational uses finitely many primes, even if the sum of all rates is infinite. A clock ring flips the sign; this is not fair resampling. In particular,

$$\mathbb{E} W_t(p) = e^{-2ta_p}.$$

For full prime noise, $a_p = 1$ and the local scale is $t \log\log X$. For selective completion, the rates tend to zero and vanish outside a reservoir adapted to P; the scale is $t\Theta_X$, where

$$\Theta_X = \sum_{p \leq X} \frac{a_p}{p} \asymp_P \log\log\log X.$$

The last scale belongs to our convenient choice of rates; it is not claimed to be an optimal or intrinsic scale for all completions. The polynomial classification in Section 6 concerns full noise and does not transfer unchanged to trace-preserving noise.

Disagreement on primes and on integers. For selective noise, Proposition 5.2 contrasts three costs on one probability-one event: relative density zero for the changed primes, asymptotic density $1/2$ for the changed integers, and infinite reciprocal sum over the changed primes at every positive time. The first limit is uniform on $[0,T]$, the second on positive compact intervals $[u,T]$; no uniform speed is asserted for the third. Thus this completion does not have finite harmonic prime budget. A sparse set of changed prime coordinates need not produce a small change in the integer sequence.

1.3 Prior work and the contribution of the transfer

The existence of root-free primes for an irreducible polynomial of degree at least two is classical: see Serre's account of Jordan's theorem and its Chebotarev application [22, Theorems 1 and 4, Section 4]. We use the harmonic form supplied by Arango-Piñeros–Keliher–Keyes [1, Theorem A]. Common reservoirs for several polynomials are governed by the Galois criteria of Berend–Bilu [3, Theorem 2] and Elsholtz–Klahn–Technau [8, Theorem 2.1 and Proposition 2.2]. Neither the reservoir nor this criterion is a new ingredient here.

There is also a direct antecedent for the constant source. Theorem 2.1 of Klurman–Mangerel–Teräväinen [14], with numbering as in arXiv:2304.05344v2, gives (1.2) for

$$w(n) = (-1)^{\Omega_E(n)}$$

whenever E has relative density zero among the primes and $\sum_{p\in E} 1/p = \infty$. Here Ω_E counts prime factors in E with multiplicity. Choosing such an E inside a root-free reservoir, beyond the prescribed prefix, already gives Theorem 1.1's existence conclusion for $f \equiv 1$. Initial translations of the index accommodate the integer constants in (1.1). A sparse set of the required type can be obtained by independently retaining each reservoir prime with probability $1/b(p)$ from Section 2.2, and fixing a realisation: the prime density follows from bounded-variable concentration, while the harmonic mass diverges because its centred series has summable variances. We do not claim this constant-source consequence as a separate novelty.

The arbitrary-source conclusion does not follow just by multiplying such a w by f: taking $f = w$ would give $fw \equiv 1$. The mechanism used here is instead Theorem 5.1, which regularises against *arbitrary bounded weights fixed before the noise*. No entropy or complexity assumption is imposed on the weights. This use of "fixed" is different from the zero-entropy meaning of "deterministic sequence" preceding Corollary 2.2 in [14]. For each prescribed source and weights, the construction yields a probability-one event; there is no single event for all sources chosen after observing the clocks.

The technical transfer is Proposition 4.2: arithmetic degree and incidence bounds are converted into a uniform-in-time estimate by revealing the small-prime trajectories, enriching a regular time mesh with their jumps, and controlling the fraction of origins affected between mesh points. Theorem 5.1 then supplies the completion, the marked law and a vanishing-delay version. The concentration inequality itself is the classical bounded-differences inequality of McDiarmid [16]; Warnke [25, Theorem 1] recalls the form used here, rather than providing an additional "typical differences" input.

For probabilistic context, Fish [9] established almost-sure normality for fair prime signs. Najnudel [18] treats consecutive values for prime laws uniform on the circle or on roots of unity; his Section 7 studies more general prime laws for one coordinate. Multiplying fair independent signs by a fixed source would not change their law, so arbitrary-source language alone does not extend that fair model. Here the source is combined with biased, non-stationary, selective noise and exact protection. Biased random multiplicative functions have also been studied by Aymone–Sidoravicius [2], in a squarefree-supported model and for partial-sum questions, rather than the completely multiplicative multivariate completion considered here.

The remaining analytic input is a root-counting Mertens formula [11]. Boolean noise and Fourier analysis [19] are classical. For full noise, the result is their quantitative transfer to a moving empirical arithmetic law, including the random translation from rational constants. All transfer arguments used below are given explicitly. No relation profile for growing words is imported.

Table 1: What is sampled and what the observer may choose. The source is always prescribed before the noise.

Observation	Prime support and scale	Sampling measure and allowed information
Prescribed trace (§5)	Root-free primes beyond the prefix, Poisson rates a_p. Local time $t\Theta_X$.	Ordinary mean over origins. The source, affine forms, marks and bounded weights precede the clocks.
Polynomial values (§6)	All primes, unit-rate Poisson clocks. At local time $\tau = t \log\log X$, two active factors give $e^{-4\tau}$.	Ordinary empirical law of fixed polynomial observations. Rational constants give a common random translation.
Prime products (§9)	Independent fair signs on the sparse set $\mathcal{D}$. One draw for all scales.	$\Lambda(r)\Lambda(s)/(X \log X)$. The joint table is prescribed. Row and column weights may follow the draw.
Band signs (§10)	One shared fair bit per band of width h in $\log\log p$. The negative class in Corollary 10.2 protects $n^2 + 1$.	Same prime-product measure and separated weights. For $h < \log 2$ dyadic tables may vary in advance. One fixed table gives all real scales.
Cofactor bit and twists (§11)	Independent flips of probability δ_p transmit $e^{-\theta/2}$ at the critical cutoff. The twists are prescribed.	Incidences $n = am = bk - 2$, counted per label pair. The bit splits each cell before taking moduli. The near-twist equivalence uses fixed odd dilations.
Periodic observations (§12)	Independent flips with probability t/H on a moving band of harmonic mass H.	Ordinary dyadic average, prescribed array. The period and bounded residue weights may follow the draw.
Fluctuations (§12.3)	Flips of probability $p/(4Z)$ on the band, or $1/4$ on one randomly selected dyadic subinterval.	Ordinary dyadic average, prescribed real array. Two independent draws on the same subinterval give the signed comparison in (12.6).

1.4 Reading the proofs

To follow the completion theorem, read Sections 2, 3.2, 4.2 and 5. Section 4.1 additionally describes its finite-local-time memory. For the polynomial channel under full noise, read Sections 3.1, 4, 6 and 7. The examples in Sections 8.1–8.2 illustrate that channel; Section 8.3 returns to selective protection.

Both routes use the same sequence of ideas. Identify primes acting to odd parity on an observation. Estimate their typical total rate. Count how many origins a single prime can affect. Convert these arithmetic budgets into temporal concentration, before applying finite Fourier analysis to the desired joint law.

For observations chosen after a draw, Section 9 removes the separated coefficients by a row-pair argument. Section 10 counts the independent bits remaining after primes are grouped into bands. Section 11 deals with a different sampling measure, in which each divisor pair labels its own progression. Its bit calculation and its comparison of nearby functions do not require the prime-pair estimates.

Table 1 records the sampling measures and the choices that may depend on the draw. For the last two rows Section 12 counts the origins changed by one prime before summing across residue classes. Appendix D uses the same divisibility counts to compute the linear part of a fluctuation.

2 Invisible primes and sparse clocks

2.1 A root-free reservoir

Lemma 2.1 (Root-free reservoir: Jordan–Chebotarev and Mertens)**.** *For P as in Theorem 1.1, there are a prime set S, disjoint from any prescribed finite set, and a number $\delta > 0$ such that*

$$p \nmid P(n) \quad (p \in S,\ n \in \mathbb{Z}), \qquad \sum_{\substack{p \le x \\ p \in S}} \frac{1}{p} = \delta \log\log x + O_P(1). \tag{2.1}$$

The remainder constant may also depend on the removed finite set.

Proof. We recall the classical argument; see [22]. Let K be the splitting field and $G = \mathrm{Gal}(K/\mathbb{Q})$ act on the $d = \deg P$ roots. This action is transitive. Burnside's formula gives mean fixed-point count one. If every element had a fixed point, the identity, with $d > 1$ fixed points, would make this mean exceed one. Thus the set of derangements is non-empty and invariant under conjugation.

Remove the primes dividing the content, leading coefficient or discriminant, the ramified primes and the prescribed finite set. At every remaining prime, the number of roots modulo p is the number of fixed points of Frobenius on the roots. The derangement classes therefore supply a root-free Chebotarev set of density

$$\delta = \frac{\#\{\sigma \in G : \sigma \text{ fixes no root}\}}{|G|} > 0.$$

The Mertens theorem for these classes [1, Theorem A], after taking the logarithm of its product and summing the p^{-2} terms, gives (2.1). A prime without a root divides no integer value of P, proving the pointwise assertion. □

The protection is exact. It uses neither a density estimate for the values of P nor a conjecture about their large prime factors.

2.2 An explicit choice of rates

The rates must serve three purposes. The condition $a_p \to 0$ makes the primes that ring on a bounded horizon sparse. Divergence of $\sum_p a_p/p$ leaves enough cumulative action for dissipation. Finiteness of $\sum_p a_p^2/p$ gives a bounded weighted second moment in Lemma 3.2. Our choice has the simple asymptotic $a_p \sim (\log 2)/\log\log p$ along S; the integer-valued denominator below makes it exact from the first prime onwards.

For a set S satisfying (2.1), put

$$b(p) = 1 + \lfloor \log_2 \lfloor \log_2 p \rfloor \rfloor, \qquad a_p = \frac{\mathbf{1}_{p \in S}}{b(p)}, \qquad \Theta_X = \sum_{p \le X} \frac{a_p}{p}. \tag{2.2}$$

This is defined from $p = 2$ onwards. It is a convenient choice, not an optimality assertion.

Lemma 2.2 (Rate budget)**.** *With $P_X = \lceil (\log X)^4 \rceil$, we have*

$$a_p \longrightarrow 0, \qquad \Theta_X = \delta(\log 2)\log\log\log X + O_S(1) \longrightarrow \infty, \tag{2.3}$$

$$\sum_p \frac{a_p^2}{p} < \infty, \qquad \Theta_{P_X} = O_S\big(\log(2 + \Theta_X)\big) = o(\Theta_X). \tag{2.4}$$

Proof. We have $b(p) = \log\log p/\log 2 + O(1)$, and hence, for sufficiently large p,

$$\frac{1}{b(p)} = \frac{\log 2}{\log\log p} + O\Big(\frac{1}{(\log\log p)^2}\Big).$$

Partial summation in (2.1) gives $\delta(\log 2)\log\log\log X + O_S(1)$ for the sum of the first term weighted by $1/p$. The error is summable, as is a_p^2/p. Small primes only change a constant. Substituting $P_X = (\log X)^4 + O(1)$ into the asymptotic proves the last assertion. □

For Theorem 1.1, all primes $p \le M_0$ are removed. The clock support S may have positive relative density. It is the set of primes whose signs actually change, or even whose clocks ring during a fixed bounded horizon, that will have relative density zero.

3 The two arithmetic budgets

Unless the dyadic alternative is explicitly selected, averages are over

$$I_X = \{n_0 + 1, \dots, n_0 + X\},$$

with n_0 fixed before the noise and large enough to make all arguments positive and non-zero. The same proofs hold for $[X, 2X) \cap \mathbb{Z}$. Implicit constants may depend on the fixed polynomials and on n_0. We always take X sufficiently large for $h = \log\log X$ to be positive.

3.1 Full noise: the number of irreducible factors

A prime matters to a sign observation precisely when its valuation is odd. The typical number of such primes will be $k\log\log X$, where k counts irreducible factors, not the algebraic degree: $T^2 + 1$ has $k = 1$, whereas $T(T + 1)$ has $k = 2$. This is why the limiting channel has one clock per factor. The calculation below is not a sieve asserting that many polynomial values are squarefree. For a primitive, squarefree, non-constant and eventually positive polynomial $R \in \mathbb{Z}[T]$, let k be its number of distinct irreducible factors over $\mathbb{Q}$, and put $D = \deg R$. If $\rho_R(p)$ counts roots modulo p, the classical input is

$$\sum_{p \le z} \frac{\rho_R(p)}{p} = k\log\log z + O_R(1). \tag{3.1}$$

This follows, for example, from Theorems 1–2 of Garcia–Lee–Suh–Yu [11]. The family is fixed: no uniformity in a varying discriminant is used. Define

$$S_R(n) = \{p : v_p(R(n)) \text{ is odd}\}, \quad K_R(n) = |S_R(n)|, \quad A_p = \#\{n \in I_X : p \in S_R(n)\}.$$

The incidence A_p is the number of origins at which prime p can change the observed sign. A flip changes a normalised average by at most $2A_p/X$, so the squared incidences in the next lemma are exactly the bounded-differences budget for large primes. The sum of incidences controls the expected fraction of origins touched in a short time interval. The centred moment ensures that few origins have too little noise. These are the three inputs of Proposition 4.2.

Lemma 3.1 (Typical degree and influences)**.** *With $P_X = \lceil (\log X)^4 \rceil$ and $h = \log\log X$,*

$$\frac{1}{X}\sum_{n \in I_X} (K_R(n) - kh)^2 \ll_R h, \tag{3.2}$$

$$\frac{1}{X^2}\sum_{n,m \in I_X} |S_R(n) \cap S_R(m)| \ll_R 1, \tag{3.3}$$

$$\sum_{p > P_X} (A_p/X)^2 \ll_R P_X^{-1}, \qquad \sum_p A_p/X \ll_R h. \tag{3.4}$$

Estimate (3.2) *remains valid with* $K_R(n)$ *replaced by* $|S_R(n)\cap(P_X,\infty)|$ *and with the same centre* kh.

Proof. First remove primes dividing the leading coefficient or discriminant. Their contribution to K_R is bounded. At all other primes the roots are simple, and their number is unchanged on lifting to modulus p^2.

Put $z = X^{1/4}$ and $V(n) = \sum^*_{p\le z} \mathbf{1}_{p|R(n)}$, where the star removes the exceptional primes. Counting congruence classes gives, for $p \neq q$,

$$\frac{1}{X}\sum_{n\in I_X} \mathbf{1}_{p|R(n)} = \frac{\rho_R(p)}{p} + O_R(X^{-1}), \qquad \frac{1}{X}\sum_{n\in I_X} \mathbf{1}_{pq|R(n)} = \frac{\rho_R(p)\rho_R(q)}{pq} + O_R(X^{-1}).$$

Expanding the square and applying (3.1) yields

$$\frac{1}{X}\sum_{n\in I_X} (V(n)-kh)^2 \ll_R h + \frac{z^2+zh}{X} \ll_R h. \tag{3.5}$$

The discrepancy between divisibility and odd valuation among these primes is bounded by

$$J(n) = \sum_{p\le z}^{*} \mathbf{1}_{p^2|R(n)}.$$

Its second moment satisfies

$$\frac{1}{X}\sum_{n\in I_X} J(n)^2 \ll_R \sum_p p^{-2} + \left(\sum_p p^{-2}\right)^2 + \frac{z^2}{X} \ll_R 1.$$

Indeed, the main terms come from congruences modulo p^2 and p^2q^2, and the $O(z^2)$ counting errors contribute $O_R(z^2/X)$. No assertion that many values $R(n)$ are squarefree is needed. Beyond z, each value $R(n) \ll_R X^D$ has only $O_R(1)$ distinct prime factors. This proves (3.2).

For the truncated version, start the sum at P_X. Its centre is

$$k\log\log z - k\log\log P_X + O_R(1) = kh + O_R(\log h).$$

The square of the displacement is $o(h)$, so the same argument applies.

Since R is primitive, it is not identically zero modulo any prime, and

$$A_p \le D(X/p+1). \tag{3.6}$$

For $P_X < p \le X$, summing the squares in (3.6) gives $O_R(P_X^{-1} + \log X/X)$. For $p > X$, there are at most D relevant indices, so

$$\sum_{p>X} \frac{A_p^2}{X^2} \le \frac{D}{X^2}\sum_p A_p \ll_R \frac{\log X}{X},$$

using the elementary bound $\omega(R(n)) \le \log R(n)/\log 2$. This handles prime divisors as large as X^D without summing a constant error over every prime up to that size. Without the cutoff, the same estimates give $\sum_p (A_p/X)^2 \ll_R 1$. The identity

$$\sum_{n,m\in I_X} |S_R(n)\cap S_R(m)| = \sum_p A_p^2$$

proves (3.3). Finally, (3.2) gives $\sum_p A_p/X = kh + O_R(\sqrt{h})$. □

Polynomial factors with even exponent disappear completely from the sign observation.

3.2 Selective noise: weighted degree of an affine product

For selective noise the same incidences are weighted by the clock rates. The typical total rate is now $k\Theta_X$; its centred second moment is bounded, thanks to $\sum_p a_p^2/p < \infty$. In the temporal argument, $\sum_p a_p A_p/X$ plays the role previously taken by the unweighted incidence sum. For an admissible family, put $R(n) = \prod_{i=1}^k \ell_i(n)$ and

$$S_R(n) = \{p : v_p(R(n)) \text{ is odd}\}, \qquad d_R(n) = \sum_{p \in S_R(n)} a_p.$$

The rates now are (2.2), and in particular vanish outside S. Keep $A_p = \#\{n \in I_X : p \in S_R(n)\}$.

Lemma 3.2 (Weighted moments)**.** *We have*

$$\frac{1}{X} \sum_{n \in I_X} (d_R(n) - k\Theta_X)^2 \ll_{\ell,S} 1, \tag{3.7}$$

$$\frac{1}{X^2} \sum_{n,m \in I_X} \sum_{p \in S_R(n) \cap S_R(m)} a_p \ll_\ell 1, \tag{3.8}$$

$$\sum_{p > P_X} (A_p/X)^2 \ll_\ell P_X^{-1}, \qquad \sum_p a_p A_p/X \ll_{\ell,S} 1 + \Theta_X. \tag{3.9}$$

After removing primes $p \le P_X$, the weighted degree still has bounded centred second moment, with centre $k(\Theta_X - \Theta_{P_X})$.

Proof. The primes dividing some A_i or some determinant $A_iB_j - A_jB_i$ form a finite set and contribute a bounded amount to the degree. At every other prime, R has exactly k simple roots modulo p and p^2, and k^2 roots modulo pq and p^2q^2 for distinct primes p, q.

With $z = X^{1/4}$ and $V(n) = \sum_{p \le z}^* a_p \mathbf{1}_{p|R(n)}$, the same congruence counts give

$$\frac{1}{X} \sum_{n \in I_X} \left(V(n) - k \sum_{p \le z}^* \frac{a_p}{p} \right)^2 \ll_\ell \sum_p \frac{a_p^2}{p} + \frac{z^2 + z(1 + \Theta_X)}{X} \ll_{\ell,S} 1.$$

The second moment of $\sum_{p \le z}^* a_p \mathbf{1}_{p^2|R(n)}$ is bounded as above. Each value $R(n) \ll_\ell X^k$ has $O_\ell(1)$ distinct prime factors exceeding z. Moreover,

$$0 \le \Theta_X - \Theta_z \le \sum_{z < p \le X} \frac{1}{p} = O(1).$$

This proves (3.7); starting at P_X proves its truncated version.

For non-exceptional primes, $A_p \le k(X/p + 1)$, and every prime dividing one of the affine forms is at most $C_\ell X$. The exceptional primes lie below P_X for large X. The sum of squared influences beyond P_X is therefore $O_\ell(P_X^{-1})$. Double counting gives

$$\frac{1}{X^2} \sum_{n,m} \sum_{p \in S_R(n) \cap S_R(m)} a_p = \sum_p a_p (A_p/X)^2 \ll_\ell 1.$$

The last assertion in (3.9) is the mean of d_R, bounded using (3.7). □

4 The common probabilistic mechanism

4.1 A finite lemma, uniform at deterministic times

Lemma 4.1 (Bias and variance of prime monomials)**.** *Let independent Poisson clocks have rates $a_p \geq 0$, let $S_1, \dots, S_M$ be finite sets, and let $|c_i| \leq 1$ be real coefficients. Put*

$$\xi_A(t) = \prod_{p \in A} (-1)^{Q_p(t)}, \quad Z(t) = \frac{1}{M} \sum_{i=1}^{M} c_i \xi_{S_i}(t), \quad d_i = \sum_{p \in S_i} a_p, \quad g_{ij} = \sum_{p \in S_i \cap S_j} a_p.$$

For $\mu > 0$, define

$$b = \frac{1}{M} \#\{i : d_i < \mu/2\}, \quad D_1 = \frac{1}{M} \sum_i |d_i - \mu|, \quad G = \frac{1}{M^2} \sum_{i,j} g_{ij}, \quad \bar{c} = \frac{1}{M} \sum_i c_i.$$

Then, for every $t \geq 0$,

$$|\mathbb{E} Z(t) - e^{-2t\mu} \bar{c}| \leq b + \frac{2D_1}{\mu}, \qquad \operatorname{Var} Z(t) \leq 2b + \frac{8G}{\mu}. \tag{4.1}$$

Proof. Independence at the primes gives the exact identities

$$\mathbb{E} \xi_{S_i}(t) = e^{-2td_i}, \tag{4.2}$$

$$\operatorname{Cov}(\xi_{S_i}(t), \xi_{S_j}(t)) = e^{-2t(d_i + d_j - 2g_{ij})}(1 - e^{-4tg_{ij}}) \geq 0. \tag{4.3}$$

When $d_i \geq \mu/2$, the mean value theorem gives

$$|e^{-2td_i} - e^{-2t\mu}| \leq 2te^{-t\mu} |d_i - \mu| \leq 2|d_i - \mu|/\mu.$$

All remaining indices cost at most b, proving the bias estimate.

For the variance, the non-negative covariances allow us to bound $c_i c_j$ in absolute value. Pairs containing a small degree cost at most $2b$. Pairs with $g_{ij} > \mu/4$ have proportion at most $4G/\mu$. For the other pairs, $d_i + d_j - 2g_{ij} \geq \mu/2$ and the covariance is at most

$$4tg_{ij}e^{-t\mu} \leq 4g_{ij}/\mu.$$

Summation gives the stated bound, including the diagonal pairs. □

The monomials need not be independent. A bound for the small-degree proportion and for mean overlap suffices. Together with Lemmas 3.1 and 3.2, the lemma gives L^2 errors $O(h^{-1/2})$ and $O(\Theta_X^{-1/2})$, respectively, uniform at deterministic times. Such estimates do not by themselves control a random temporal supremum.

4.2 Concentration on an entire time interval

Small primes can affect many origins at once. We therefore reveal their entire paths: they remain present as a known time-dependent mask, rather than being discarded or frozen. Large primes have small individual influence, so their squared-incidence budget gives conditional concentration at a fixed time, including a time determined only by the small-prime paths.

To pass between such times, enrich a regular mesh with every small-prime jump. Within a mesh cell, control the fraction of origins touched by a large-prime ring, not the number of those rings. An origin touched repeatedly still contributes at most a difference of two signs. This is the step that turns fixed-time concentration into a bound for a whole trajectory.

Proposition 4.2 (Temporal concentration after revealing small-prime paths)**.** *For each integer X, let $S_X(n)$, $n \in I_X$, be finite prime supports, and put $P_X = \lceil (\log X)^4 \rceil$. Let the clock rates satisfy $0 \le a_p \le 1$, and let $|c_n| \le 1$ be arbitrary real coefficients fixed before the noise. Write*

$$A_p = \#\{n : p \in S_X(n)\}, \quad d_{>P_X}(n) = \sum_{\substack{p \in S_X(n) \\ p > P_X}} a_p, \quad Z_X(t) = \frac{1}{X} \sum_{n \in I_X} c_n \xi_{S_X(n)}(t).$$

Assume $\Lambda_X \to \infty$, $\Lambda_X \le \log X$ for all sufficiently large X, and

$$\frac{1}{X} \#\{n : d_{>P_X}(n) < \Lambda_X/2\} \le C\Lambda_X^{-1}, \tag{4.4}$$

$$\sum_{p > P_X} (A_p/X)^2 \le C/P_X, \qquad \sum_p a_p A_p/X \le C\Lambda_X. \tag{4.5}$$

For each $0 < T < \infty$ there exist $c, C_T > 0$ such that, for $0 \le u \le T$ and $\varepsilon > 0$,

$$\mathbb{P}\left\{ \sup_{u \le t \le T} |Z_X(t)| > e^{-u\Lambda_X} + C_T\Lambda_X^{-1} + 3\varepsilon \right\} \le C_T(1 + P_X + \Lambda_X^2) e^{-c\varepsilon^2 P_X}. \tag{4.6}$$

Consequently, for every positive sequence u_X, fixed before the noise, with $u_X \to 0$ and $u_X\Lambda_X \to \infty$, almost surely for all sufficiently large X,

$$\sup_{u_X \le t \le T} |Z_X(t)| \ll_T e^{-u_X\Lambda_X} + \Lambda_X^{-1}. \tag{4.7}$$

The constants depend on the constants in the assumptions, but not on the coefficients or on u. The coefficients may vary with X provided they are specified before the clocks.

Proof. For each X, the union $\bigcup_{n \in I_X} S_X(n)$ is finite. Only these clocks influence Z_X, so its relevant paths have finitely many jumps on a bounded horizon, almost surely. All temporal suprema below are therefore well defined; the argument is not an intersection over uncountably many time-indexed events.

Condition on the entire trajectories of the primes $p \le P_X$. They give a known, time-dependent sign mask; their values are not frozen in time. For $t \ge u$, assumption (4.4) gives

$$|\mathbb{E}(Z_X(t) \mid \text{small-prime paths})| \le \frac{1}{X} \sum_{n \in I_X} e^{-2ud_{>P_X}(n)} \le e^{-u\Lambda_X} + C\Lambda_X^{-1}. \tag{4.8}$$

Changing one large-prime sign affects $Z_X(t)$ by at most $2A_p/X$. The classical bounded-differences inequality [16, 25], using (4.5), yields

$$\mathbb{P}\{|Z_X(t)| > e^{-u\Lambda_X} + C\Lambda_X^{-1} + \varepsilon \mid \text{small-prime paths}\} \le 2e^{-c\varepsilon^2 P_X}. \tag{4.9}$$

The estimate also holds for a time determined solely by the small-prime paths, and for the mask immediately before or after one of their rings. At a time fixed by these paths, large clocks almost surely do not ring exactly at that time.

Choose the regular deterministic mesh

$$N_X = \max\{1, \lceil T\Lambda_X^2 \rceil\}, \qquad t_j = \frac{jT}{N_X} \quad (0 \le j \le N_X).$$

It has $O_T(1 + \Lambda_X^2)$ points and step at most Λ_X^{-2}. Add u and all ring times of primes $p \le P_X$ in $[0, T]$. The enriched mesh has expected cardinality $O_T(1 + \Lambda_X^2 + P_X)$, since the rates are at most one. Conditional on the small-prime paths, these points are fixed independently of

the large clocks, so (4.9) applies at each point in $[u, T]$, recording left and right values where needed.

For one deterministic mesh cell I, let J_I be the fraction of origins whose support contains at least one large prime ringing in I. A union bound gives

$$\mathbb{E}J_I \le |I| \sum_p a_p A_p / X \le C\Lambda_X^{-1}.$$

The large-prime ringing indicators are independent. Changing one indicator affects J_I by at most A_p/X. A second application of bounded differences gives

$$\mathbb{P}\{J_I > C\Lambda_X^{-1} + \varepsilon\} \le e^{-c\varepsilon^2 P_X}.$$

Between consecutive enriched mesh points, the small-prime mask is constant and any difference of two values of Z_X is bounded by $2J_I$ in the containing deterministic cell. Multiple rings at one origin do not invalidate this bound: it controls a difference of values, not the total variation of the entire path. A conditional union bound over enriched points, a union bound over cells, and integration prove (4.6). The cost is the mass of affected origins, not the number of large-prime rings.

Finally, take $\varepsilon = 1/\log X$ and $u = u_X$. Since $P_X \asymp (\log X)^4$ and $\Lambda_X \le \log X$, the right-hand side is

$$O_T\big((\log X)^4 e^{-c(\log X)^2}\big),$$

summable over all positive integers X. Borel–Cantelli gives (4.7), since $1/\log X \le 1/\Lambda_X$. No monotonicity of u_X is needed. □

For full noise, Lemma 3.1 verifies the assumptions with $\Lambda_X = h$. For selective noise on an affine product, the truncated centre is $k(\Theta_X - \Theta_{P_X}) \sim k\Theta_X$ and the centred second moment is bounded. The proportion below $\Theta_X/2$ is therefore even $O_{\ell,S}(\Theta_X^{-2})$. Proposition 4.2 applies with $\Lambda_X = \Theta_X$.

5 Selective completion: proof and cost

5.1 Weighted cancellation and local memory

Theorem 5.1 (Affine regularisation against fixed bounded weights)**.** *Fix S satisfying* (2.1)*, rates* (2.2)*, a source $f \in \mathcal{M}_\pm$ and arbitrary real weights $|v_n| \le 1$, all before the clocks. For an admissible affine family and a positive sequence $u_X \to 0$ fixed before the noise with $u_X\Theta_X \to \infty$, almost surely,*

$$\sup_{u_X \le t \le T} \left| \frac{1}{X} \sum_{n \in I_X} v_n \prod_{i=1}^{k} F_t^f(\ell_i(n)) \right| \ll_{\ell,S,T} e^{-u_X\Theta_X} + \Theta_X^{-1} \tag{5.1}$$

for all sufficiently large X and every fixed $T < \infty$. The event may depend on f, v and u_X. No complexity or entropy assumption is made on the weights.

Proof. Write the product as $f(R(n))\xi_{S_R(n)}(t)$, where $R = \prod_i \ell_i$. Apply Proposition 4.2 with coefficients $c_n = v_n f(R(n))$, supports and estimates from Lemma 3.2, and $\Lambda_X = \Theta_X$. Intersect over positive integer horizons to obtain every finite horizon. □

The information retained at finite local time is also useful. For arbitrary real coefficients $|c_n| \le 1$ fixed before the noise, Lemmas 3.2 and 4.1, with $\mu = k\Theta_X$, give

$$\sup_{t \ge 0} \left\| \frac{1}{X} \sum_{n \in I_X} c_n W_t(R(n)) - e^{-2kt\Theta_X} \frac{1}{X} \sum_{n \in I_X} c_n \right\|_{L^2(\mathbb{P})} \ll_{\ell,S} \Theta_X^{-1/2}. \tag{5.2}$$

Indeed, $b \ll \Theta_X^{-2}$, $D_1 \ll 1$ and $G \ll 1$ in (4.1). The supremum lies outside the norm. At time $t = s/\Theta_X$, the memory multiplier is the non-zero number e^{-2ks}.

There is no conflict between this memory and regularity at every positive time. Each fixed integer prefix involves finitely many clocks, so almost surely it stays equal to its initial value for a short random interval after zero. Regularisation concerns the limit $X \to \infty$ at positive times, or after delays satisfying $u_X\Theta_X \to \infty$. These assertions do not exchange the limits in time and observation size.

5.2 Proof of the main theorem and the marked version

Proof of Theorem 1.1. Choose the reservoir in Lemma 2.1, removing all $p \le M_0$, and use rates (2.2). No active prime divides a value of P, so (1.3) holds pointwise along every trajectory at every time. Integers $n \le M_0$ involve only inactive primes, giving (1.4).

For large X, take $u_X = \Theta_X^{-1/2}$. Theorem 5.1 with $v_n = 1$, followed by a countable intersection over integer-coefficient affine families and integer horizons, gives $F_t^f \in \mathcal{L}_\pm$ at every real $t > 0$ on one probability-one event. Each such time eventually lies above u_X. Proposition 5.2 below gives the prime-density assertion, uniformly on bounded horizons. Fixing $t = 1$ on an admissible realisation supplies g. □

Proof of Corollary 1.2. For $x \in \{-1, 1\}^k$ and a mark state z, expand

$$\mathbf{1}_{(F_t^f(\ell_i(n)))_i = x}\mathbf{1}_{z(n)=z} = 2^{-k}\mathbf{1}_{z(n)=z}\prod_{i=1}^{k}(1 + x_i F_t^f(\ell_i(n))).$$

After averaging, the empty term is exactly $2^{-k}\nu_X\{z\}$. Each non-empty term is controlled by (5.1), using the fixed weight $\mathbf{1}_{z(n)=z}$ and the corresponding affine subfamily. Summing over the finitely many states gives (1.7). This also gives the vanishing-delay version with error $O(e^{-u_X\Theta_X} + \Theta_X^{-1})$ for fixed forms and marks. The arithmetic estimates and the concentration argument apply separately to the fixed-prefix and dyadic choices of I_X. □

5.3 Prime density, integer disagreement and harmonic mass

Proposition 5.2 (Prime cost and integer disagreement)**.** *For selective noise, almost surely, for every $T < \infty$,*

$$\sup_{0 \le t \le T} \frac{\#\{p \le x : F_t^f(p) \ne f(p)\}}{\pi(x)} \longrightarrow 0. \tag{5.3}$$

On one probability-one event, for every $0 < u < T < \infty$,

$$\sup_{u \le t \le T} \left| \frac{1}{X}\#\{n \le X : F_t^f(n) \ne f(n)\} - \frac{1}{2} \right| \longrightarrow 0. \tag{5.4}$$

On that event, for all real times $t > 0$ simultaneously,

$$\sum_{p : F_t^f(p) \ne f(p)} \frac{1}{p} = \infty. \tag{5.5}$$

The last assertion does not claim a uniform rate of divergence on time intervals.

Proof. For fixed T, the variables $I_p = \mathbf{1}_{Q_p(T) \ge 1}$ are independent with means $1 - e^{-Ta_p} \to 0$. Hoeffding's bounded-variable inequality, followed by Borel–Cantelli along the first m primes,

gives $m^{-1}\sum_{j\le m} I_{p_j} \to 0$. Every prime changed before time T has rung, proving (5.3). Positive integer horizons form a countable family.

The quotient $F_t^f(n)/f(n)$ equals $W_t(n)$, independently of f. Hence

$$\frac{1}{X}\#\{n \le X : F_t^f(n) \neq f(n)\} = \frac{1}{2}\left(1 - \frac{1}{X}\sum_{n\le X} W_t(n)\right).$$

Theorem 5.1 applied to the constant source and the single form n proves (5.4), on one event for all positive compact time intervals.

For the simultaneous conclusion (5.5), we use a deterministic fact. If $w \in \mathcal{M}_\pm$ and $E = \{p : w(p) = -1\}$ satisfies $\sum_{p\in E} 1/p < \infty$, then

$$\lim_{X\to\infty} \frac{1}{X}\sum_{n\le X} w(n) = \prod_{p\in E} \frac{p-1}{p+1} > 0. \tag{5.6}$$

To prove it, let w_y retain only the negative prime signs in $E \cap [2, y]$. For this finite set, truncate the valuations and apply the Chinese remainder theorem. Passing through the valuation truncations gives

$$\lim_{X\to\infty} \frac{1}{X}\sum_{n\le X} w_y(n) = \prod_{\substack{p\in E\\ p\le y}} \left((1-p^{-1})\sum_{j\ge 0}\frac{(-1)^j}{p^j}\right) = \prod_{\substack{p\in E\\ p\le y}} \frac{p-1}{p+1}.$$

The discarded valuation tails have density tending to zero, which justifies the finite-set limit. If $w(n) \neq w_y(n)$, some prime of E greater than y divides n. Uniformly in X,

$$\frac{1}{X}\sum_{n\le X} |w(n) - w_y(n)| \le 2\sum_{\substack{p\in E\\ p>y}} \frac{1}{p}. \tag{5.7}$$

This tends to zero as $y \to \infty$. The infinite product converges to a positive number because its logarithmic terms are $O(1/p)$ in absolute value for large p. This proves (5.6).

On the event already used for (5.4), the mean of W_t is zero for every $t > 0$. Apply the deterministic implication (5.6) to each of these functions W_t. Its negative-prime set cannot have finite harmonic mass. Since $W_t(p) = -1$ precisely when $F_t^f(p) \neq f(p)$, (5.5) follows *without any additional probabilistic intersection over times.*

A separate fixed-time calculation retains quantitative information about the expected truncated mass. For fixed $t > 0$, the modification indicators are independent with probabilities $q_p(t) = (1 - e^{-2ta_p})/2$, and (2.4) gives

$$\mathbb{E}\sum_{\substack{p\le X\\ F_t^f(p)\neq f(p)}} \frac{1}{p} = \sum_{p\le X}\frac{q_p(t)}{p} = t\Theta_X + O_t(1). \tag{5.8}$$

The centred series $\sum_p (\mathbf{1}_{Q_p(t)\text{ odd}} - q_p(t))/p$ converges almost surely, by summability of its variances. This is a useful fixed-time description, but the all-time assertion above comes from the deterministic argument. □

The completion is neither a negligible Hamming perturbation of the integers nor a surgery with finite harmonic prime budget. Relative prime density alone does not measure its effect on observations.

5.4 An entire family of protected norms

Root-freeness is sufficient for protection, but stronger than necessary. A prime whose valuation is always even on the protected arguments is equally invisible to a sign flip; this protects the whole family below.

Corollary 5.3 (Sums of two squares)**.** *For every source $f \in \mathcal{M}_{\pm}$ and every prescribed prefix, the construction can satisfy, at all times,*

$$F_t^f(a^2+b^2) = f(a^2+b^2) \qquad (a,b \in \mathbb{Z},\ a^2+b^2>0),$$

while retaining all the preceding completion and cost conclusions.

Proof. Take $S = \{p > M_0 : p \equiv 3 \pmod 4\}$, with any further finite removals. It satisfies (2.1) with $\delta = 1/2$. If such a p divides a^2+b^2, then it divides both a and b. Dividing by p^2 and repeating shows that $v_p(a^2+b^2)$ is even. Any sign change at these primes therefore leaves the norm value unchanged. Affine regularisation uses only the harmonic reservoir estimate and the rates. □

For $f = \lambda$, all the protected values are exactly Liouville values. For $f \equiv 1$, the completion is constantly $+1$ on these norms while belonging to $\mathcal{L}_{\pm}$. As explained above, the latter existence case follows already from [14, Theorem 2.1, arXiv v2]. It illustrates a previously available separation: affine cancellation alone does not force a vanishing polynomial mean.

6 The irreducible-factor channel

Throughout this section every prime clock has rate one. The model remains (1.8), but it no longer generally preserves the trace singled out by selective completion.

The need for a factor channel is already visible in a small example. Put $A = T^2+1$ and $B = T^2+2$. The observations (A, B, AB, A^2) always give a sign vector of the form $(x, y, xy, 1)$, not four free signs. A flip in the A factor changes the first and third coordinates; a flip in B changes the second and third; the square does not change. The matrix below records exactly these shared actions. Section 8.1 gives the four states and their limiting frequencies.

6.1 Observation algebra and the rational-constant environment

Fix $R_1, \dots, R_m \in \mathbb{Q}(T)^{\times}$, positive at all sufficiently large integers. Factor them as

$$R_j(T) = c_j \prod_{\alpha=1}^{r} P_\alpha(T)^{e_{j\alpha}}, \qquad c_j \in \mathbb{Q}_{>0}, \quad e_{j\alpha} \in \mathbb{Z}, \tag{6.1}$$

where the $P_\alpha \in \mathbb{Z}[T]$ are primitive, irreducible, distinct and have positive leading coefficient. Negative exponents are allowed. In the parity matrix, row j describes observation R_j, while column α records which observations change when the sign of factor P_α flips. Put

$$E = (e_{j\alpha} \bmod 2) \in \mathbb{F}_2^{m\times r}, \qquad H = \{(-1)^{Ez} : z \in \mathbb{F}_2^r\}, \qquad d(u) = \mathrm{wt}(E^T u) \quad (u \in \mathbb{F}_2^m). \tag{6.2}$$

Sign vectors are multiplied coordinatewise. The initial and noisy empirical laws are

$$\mu_X^f = \frac{1}{X}\sum_{n\in I_X} \delta_{(f(R_j(n)))_j}, \qquad \widehat{\mu}_{X,t}^f = \frac{1}{X}\sum_{n\in I_X} \delta_{(F_t^f(R_j(n)))_j}.$$

The interval convention of Section 3 excludes zeros, poles and negative values. The initial law may vary without converging.

Let $B_\alpha(s) = (-1)^{N_\alpha(s)}$ be independent rate-one telegraph processes starting at $+1$. The channel K_s multiplies an initial sign vector by

$$\left(\prod_\alpha B_\alpha(s)^{e_{j\alpha}}\right)_{j\le m}.$$

Writing v_α for column α of E, its generator and Fourier multipliers are

$$\mathcal{A}\varphi(x) = \sum_{\alpha=1}^{r}[\varphi(x\cdot(-1)^{v_\alpha}) - \varphi(x)], \qquad \widehat{K_s}(u) = e^{-2sd(u)}. \tag{6.3}$$

This is an image of classical Boolean noise [19, Section 2.4, Definition 2.46 and Proposition 2.47], with noise parameter e^{-2s}. Repeated columns retain their multiplicities; zero columns have no visible effect. One irreducible factor, rather than one output observation, supplies each elementary clock. The rank of E determines the size of the compatible group H, while its columns with multiplicity determine the finite-time transition law, as Section 8.2 illustrates. For $a \in \{-1,1\}^m$, write $\tau_a\mu$ for the law of aV when V has law μ.

Rational constants carry a different kind of information. For the pair $(n, 2n)$, the product of the observed signs is exactly $F_t^f(2)$. This pair is not admissible in (1.1), because its forms are proportional; it belongs to the rational-observation theory here. Its two-state compatible class moves with the clock at 2, which explains the translation in the next theorem.

Theorem 6.1 (Empirical channel with rational constants)**.** *For the fixed family* (6.1) *and* $h = \log\log X$*,*

$$\sup_{t\ge 0}\sup_{f\in\mathcal{M}_\pm}\left\|d_{\mathrm{TV}}\left(\widehat{\mu}^f_{X,t}, \tau_{(W_t(c_j))_j}(\mu_X^f K_{th})\right)\right\|_{L^2(\mathbb{P})} \ll_R h^{-1/2}. \tag{6.4}$$

Both suprema are outside the expectation. The source may depend on X as long as it is fixed before the clocks.

The target is random through primes dividing the numerators or denominators of the constants c_j. It is not assumed independent of the field. At time $t = s/h$ with bounded s, the translation equals the identity with probability $1 - O_R(s/h)$. At a fixed positive physical time, it cannot generally be dropped. "Constant environment" refers to constant rational *arguments*; their signs evolve in time.

Proof. For $u \in \mathbb{F}_2^m$, write

$$\prod_j R_j(T)^{u_j} = C_u R_u(T) A_u(T)^2, \qquad C_u = \prod_j c_j^{u_j}, \quad R_u = \prod_{\alpha:(E^Tu)_\alpha=1} P_\alpha.$$

Set $C^f_{u,X}(t) = X^{-1}\sum_{n\in I_X}\prod_j F_t^f(R_j(n))^{u_j}$. If $d(u) > 0$, Lemma 3.1 for R_u gives, in Lemma 4.1 with $\mu = d(u)h$,

$$b \ll_R h^{-1}, \qquad D_1 \ll_R \sqrt{h}, \qquad G \ll_R 1.$$

Consequently,

$$\left\|C^f_{u,X}(t) - W_t(C_u)e^{-2d(u)th}C^f_{u,X}(0)\right\|_{L^2(\mathbb{P})} \ll_R h^{-1/2}. \tag{6.5}$$

Indeed, first apply the lemma with the fixed coefficient $f(R_u(n))$, then multiply the error by $F_t^f(C_u)$. Multiplication by this random sign preserves the L^2 norm exactly, whether or not it depends on the other clocks.

When $d(u) = 0$, there is instead the exact identity

$$C^f_{u,X}(t) = F_t^f(C_u). \tag{6.6}$$

The Fourier coefficients of the target in (6.4) are precisely the target coefficients in (6.5). Finite Fourier inversion and the triangle inequality in L^2 complete the proof, with constants depending on the fixed family. □

6.2 Simultaneous classification at positive times

Theorem 6.2 (Square-compatible equidistribution)**.** *Fix $f \in \mathcal{M}_{\pm}$ and a positive sequence $u_X \to 0$, before the clocks, with $u_X \log\log X \to \infty$. Almost surely, for every fixed family* (6.1) *and every $T < \infty$,*

$$\sup_{u_X \le t \le T} d_{\mathrm{TV}}\Big(\widehat{\mu}^f_{X,t}, \mathrm{Unif}(F^f_t(c)H)\Big) \ll_{R,T} e^{-u_X h} + h^{-1} \tag{6.7}$$

for all sufficiently large X, where $F^f_t(c) = (F^f_t(c_j))_{j\le m}$. A single realisation satisfies the resulting classification for all fixed rational families and all real $t > 0$.

Proof. For every u with $d(u) > 0$, apply Proposition 4.2 to R_u, with $\Lambda_X = h$ and $c_n = f(R_u(n))$. Lemma 3.1 verifies its assumptions. Multiplication by $F^f_t(C_u)$ does not change absolute values, even at a ring affecting a constant. These characters tend uniformly to zero at the asserted rate.

The characters with $E^T u = 0$ are exactly those annihilating H. By (6.6), they equal $F^f_t(C_u)$ at every size and time. These are the Fourier coefficients of the uniform law on $F^f_t(c)H$; all its other coefficients vanish. Fourier inversion proves (6.7). Rational families and integer horizons are countable, so we can intersect their probability-one events. The event depends on the fixed source, and each $t > 0$ is eventually above u_X. □

Square identities are thus the exact invariant characters of the channel. The initial law is supported *within* $f(c)H$; not every point of that coset need occur at every finite size. The theorem describes the frequencies after regularisation.

Corollary 6.3 (Normality along a polynomial)**.** *Let $P \in \mathbb{Q}[T]$ be eventually positive. If it is not of the form $cQ_0(T)^2$ with $c \in \mathbb{Q}_{>0}$ and $Q_0 \in \mathbb{Q}[T]$, then for every source f fixed before the clocks, almost surely*

$$n \longmapsto F^f_t(P(n))$$

is normal for all real times $t > 0$ simultaneously. Word frequencies converge uniformly on compact time intervals contained in $(0,\infty)$, with the vanishing-delay version in (6.7)*. When $P = cQ_0^2$, the sequence is instead constantly $F^f_t(c)$ outside the omitted initial segment.*

Proof. Take a non-empty set of distinct integer shifts A. Among the odd-multiplicity roots of P, choose one, α, with maximal real part, and put $a_0 = \min A$. The root $\alpha - a_0$ occurs to odd multiplicity in $P(T+a_0)$. If it occurred to odd multiplicity in another $P(T+a)$, then $\alpha+(a-a_0)$ would be an odd-multiplicity root of P with strictly greater real part. Any even-multiplicity occurrences do not change the parity. Hence $\prod_{a\in A} P(T+a)$ is not a constant times a square. All non-trivial word characters vanish by Theorem 6.2, giving frequencies $2^{-|A|}$. The square case follows directly from complete multiplicativity. □

This corollary does not follow merely from normality on consecutive integers; the theorem is applied to the polynomial arguments themselves. It is also a first-order statement at scale X, not a central limit theorem at scale $\sqrt{X}$. For comparison, de la Bretèche–Wang–Xu [5, Theorem 1.2] prove a Gaussian limit for the fair completely multiplicative model, called extended Rademacher, when $P \in \mathbb{Z}[T]$ is separable, has positive leading coefficient and total degree at least two, and all its irreducible factors have degree at most three. Our conclusion requires less asymptotic precision and concerns a different, non-stationary prime law. It is not an extension of that central limit theorem to all degrees.

7 Local paths and memory of the source

There are two distinct path questions. Proposition 7.1 compares an entire path at one sampled origin after integrating both the origin and the clocks. Proposition 7.2 instead compares an empirical law at finitely many times after fixing the clocks, with its error averaged over the clocks. The second statement needs an additional variance estimate; a full-path empirical comparison in total variation is obstructed by finite support, as explained below.

7.1 Coupling at a random origin

The noise is still full, with rate one at every prime. Fix $S_0 > 0$ and choose U_X uniformly in I_X, independently of the clocks. For $0 \le s \le S_0$, observe

$$\gamma_X(s) = \big(F^f_{s/h}(R_j(U_X))\big)_{j\le m}.$$

The target γ°_X starts from the same vector $(f(R_j(U_X)))_j$ and multiplies it by the factor telegraphs in (6.3).

Proposition 7.1 (Local path comparison, averaged over the clocks)**.** *On the space of càdlàg paths, retaining the jump times,*

$$d_{\mathrm{TV}}\big(\mathcal{L}(U_X,\gamma_X),\mathcal{L}(U_X,\gamma^\circ_X)\big) \ll_R S_0 h^{-1/2}. \tag{7.1}$$

Both the origin and the clocks are integrated in this comparison.

Proof. Remove a finite set B containing the primes in the constants and the resultants between distinct irreducible factors. Outside B, the prime supports of the values $P_\alpha(n)$ are disjoint at a fixed origin: a common prime would divide their non-zero resultant. The probability that any prime of B rings before S_0/h is at most $S_0|B|/h$.

At a fixed origin n, on the event of no such ring, the products of signs outside B are independent telegraphs of rates $K_{P_\alpha,\notin B}(n)/h$ in local time. Coupling rate-r and rate-one Poisson processes through their common part costs at most $S_0|r-1|$ on the horizon. Lemma 3.1 and Cauchy–Schwarz give

$$\frac{1}{X}\sum_{n\in I_X} \big|K_{P_\alpha,\notin B}(n)/h - 1\big| \ll_R h^{-1/2}.$$

Sum over the factors and average the origins. This proves (7.1). The coupling treats one sampled origin, not all positions simultaneously. □

7.2 Finite-time empirical laws after fixing the clocks

Proposition 7.2 (Empirical finite-time marginals)**.** *Fix distinct integer shifts $a_1,\dots,a_m$, a source $f \in \mathcal{M}_\pm$ before the clocks, and local times $s_1,\dots,s_q \in [0,S_0]$. Use either interval convention from Section 3, with all shifted arguments positive. Define the random empirical law*

$$\widehat{\nu}_{X,\mathbf{s}}(\omega) = \frac{1}{X}\sum_{n\in I_X} \delta_{(F^f_{s_\ell/h}(n+a_i;\omega))_{i,\ell}} \tag{7.2}$$

and the deterministic target

$$\nu^\circ_{X,\mathbf{s}} = \frac{1}{X}\sum_{n\in I_X} \mathcal{L}\big((f(n+a_i)B_i(s_\ell))_{i,\ell}\big), \tag{7.3}$$

where $B_1, \dots, B_m$ are independent rate-one telegraphs starting at $+1$. Then

$$\mathbb{E}_\omega d_{\mathrm{TV}}(\widehat{\nu}_{X,\mathbf{s}}(\omega), \nu^\circ_{X,\mathbf{s}}) \ll_{a,q,S_0} h^{-1/2}. \tag{7.4}$$

For each coordinate i, the same telegraph is observed at all the times s_ℓ; fresh independent signs are not drawn at successive times.

Proof. This needs a variance estimate beyond the coupling in Proposition 7.1. For a fixed test $0 \le \Phi \le 1$ of the local path, replace one whole prime-clock history on $[0, S_0/h]$ by an independent copy. This changes its empirical average over origins by at most

$$b_p \le m(1/p + 1/X).$$

Only primes $p \le C_a X$ can occur in the shifted arguments, for either interval convention, and $\sum_p b_p^2 \ll_a 1$ when restricted to these primes. The two histories are both empty with probability $e^{-2S_0/h}$, in which case the replacement changes nothing. The Efron–Stein product variance inequality [6, 24] yields

$$\mathrm{Var}\left(\frac{1}{X}\sum_{n\in I_X}\Phi(\gamma_{X,n})\right) \le \frac{1}{2}(1-e^{-2S_0/h})\sum_p b_p^2 \ll_a S_0/h. \tag{7.5}$$

Apply this to each indicator of a state of the mq sampled signs. Cauchy–Schwarz and summation over the fixed finite state space bound $\mathbb{E}d_{\mathrm{TV}}(\widehat{\nu}_{X,\mathbf{s}}, \mathbb{E}\widehat{\nu}_{X,\mathbf{s}})$ by $O_{a,q,S_0}(h^{-1/2})$. Projecting the path coupling (7.1) onto the selected times bounds $d_{\mathrm{TV}}(\mathbb{E}\widehat{\nu}_{X,\mathbf{s}}, \nu^\circ_{X,\mathbf{s}})$ by the same order. The triangle inequality proves (7.4). □

This assertion retains finitely many times and averages its error over the clocks. It is neither an almost-sure path statement nor total-variation convergence of empirical laws of entire paths. Indeed, once the clocks are fixed, such an empirical path law is supported on at most X paths. For a general factor matrix E, let r_* be the number of its non-zero columns, counted with multiplicity. The target has a nonatomic component of mass $1 - e^{-r_* S_0}$, corresponding to at least one visible jump at a continuous time. Consequently,

$$d_{\mathrm{TV}}\big(\text{empirical path law}, \mathcal{L}(\gamma^\circ_X)\big) \ge 1 - e^{-r_* S_0}. \tag{7.6}$$

Repeated non-zero columns do not remove this component: their independent clocks almost surely have distinct jump times. This is a finite-support obstruction, not an additional arithmetic cancellation claim.

7.3 The universal scale and the limitation of returning to time zero

For distinct shifts alone, the irreducible factors are the corresponding affine polynomials and the channel is

$$B_s(v, w) = \prod_{j=1}^{m} \frac{1 + e^{-2s} v_j w_j}{2}.$$

Theorem 6.1 gives the following regimes, in probability and for a fixed family size.

Scale	Empirical law after noise, up to $o(1)$
$t_X h \to 0$	Moving initial law μ_X^f.
$t_X h \to s \in (0, \infty)$	$\mu_X^f B_s$, with non-zero memory.
$t_X h \to \infty$	Uniform law on the cube.

The last condition is necessary for a guarantee valid for every source: for $f \equiv 1$, at finite local time s, a single-letter mean tends to $e^{-2s} > 0$. If $t_X h$ fails to tend to infinity, one can pass to a subsequence converging to a finite s. This necessity is not asserted for each source individually; a source may already be normal before noise.

For a non-empty set of shifts A and the Liouville source, (6.5) at time s/h reads

$$\frac{1}{X} \sum_{n \in I_X} \prod_{a \in A} F^{\lambda}_{s/h}(n+a) = e^{-2|A|s} \frac{1}{X} \sum_{n \in I_X} \prod_{a \in A} \lambda(n+a) + O_{L^2}(h^{-1/2}). \tag{7.7}$$

At fixed s, convergence of the left-hand side to zero in probability would therefore be equivalent to cancellation of the deterministic correlation on the right. The theorem calculates the multiplier, not this correlation. When $t_X h \to \infty$, memory disappears; cancellation of the output cannot then be inverted without additional information. The same distinction holds for selective noise, with Θ_X and (5.2).

Finally, the order of choices is indispensable. Choosing $f(p) = (-1)^{Q_p(t_0)}$ after observing the clocks would give $F^f_{t_0} \equiv 1$. There can be no single event asserting these conclusions for all adaptive sources simultaneously.

8 Channel examples and the limits of protection

8.1 Four observations need not be four free signs

Take $A = T^2 + 1$, $B = T^2 + 2$, and observations (A, B, AB, A^2). Their parity matrix is

$$E = \begin{pmatrix} 1 & 0 \\ 0 & 1 \\ 1 & 1 \\ 0 & 0 \end{pmatrix}.$$

The only permitted states are

$$(+,+,+,+), \quad (-,+,-,+), \quad (+,-,-,+), \quad (-,-,+,+).$$

Theorem 6.2 gives each asymptotic frequency $1/4$. At every size and time their law has total-variation distance at least $3/4$ from four independent fair signs, since the other twelve states are forbidden.

For $(n, 2n)$, the constant environment is essential: the product of the two signs equals $F^f_t(2)$ exactly. The compatible limiting law is uniform on

$$\{(1, F^f_t(2)),\ (-1, -F^f_t(2))\}.$$

8.2 Column multiplicities matter

For $P(T) = T(T+1)$, the observations $P(n)$ and $P(n+1)$ have columns

$$(1,0)^T, \qquad (1,1)^T, \qquad (0,1)^T.$$

Each non-trivial character has $d(u) = 2$. From the constant source, at local time with $e^{-2s} = 1/2$, the probabilities in the order $(++,\ -+,\ +-,\ --)$ are

$$(7,3,3,3)/16.$$

They are not the probabilities of two independent clocks acting directly on the outputs. As time tends to infinity, rank two does give four frequencies equal to 1/4.

More generally, if m_v is the multiplicity of a non-zero column v, then

$$d(u) = \sum_{v \neq 0} m_v \frac{1 - (-1)^{\langle u,v \rangle}}{2}, \qquad m_v = -2^{1-m} \sum_{u \in \mathbb{F}_2^m} d(u)(-1)^{\langle u,v \rangle}.$$

This classical Fourier inversion recovers the visible incidences of the channel. It does not recover zero columns and is not a new polynomial factorisation algorithm.

8.3 Several traces: a sufficient criterion and its specific obstruction

The availability of a common root-free reservoir is a classical Galois question, not a new criterion of this paper. The density description in Berend–Bilu [3, Theorem 2], as restated in Elsholtz–Klahn–Technau [8, Theorem 2.1 and Proposition 2.2], identifies primes where a polynomial has a root by fixed points in the Galois action. The following proposition applies that criterion to our completion construction.

Proposition 8.1 (Completion from a common reservoir)**.** *Let $P_1, \dots, P_r$ be non-constant, separable, eventually positive integer polynomials. If the Galois group of their common splitting field has an element fixing no point in the union of their roots, the conclusions of Theorem 1.1 and Corollary 1.2 hold while preserving all the traces P_j simultaneously.*

Proof. The conjugacy classes with no fixed point in this union provide a positive-density Chebotarev set root-free for every P_j. After finite removals, the harmonic conclusion of Lemma 2.1 and the whole weighted construction apply without further change. □

The examples below give the relative density of the common root-free reservoir, outside exceptional primes.

Protected traces	Reservoir density
$T^2 + 1$	1/2
$T^3 - T - 1$	1/3
$T^2 - 2$ and $T^2 - 3$	1/4
$T^2 - 2$, $T^2 - 3$ and $T^2 - 6$	0

The exact product in the last row appears in [8, equation (1.1)]. It is *exceptional*, meaning that it has a root modulo all but finitely many primes and no integer root, but is not *intersective*: it has no root modulo 2^6. Root availability modulo almost all primes is the property relevant here, not root availability modulo every integer.

The cubic is irreducible by the rational-root test, and its discriminant -23 gives Galois group S_3, with two derangements out of six elements. The two independent quadratics have group $\mathbb{F}_2^2$. For the last row, all three quadratic symbols cannot be negative simultaneously, since $(6/p) = (2/p)(3/p)$ for $p > 3$.

Absence of a reservoir has a specific, limited consequence. Outside a finite exceptional set, a simple root of P_j modulo p has exactly one root lift among its p lifts modulo p^2. The other lifts provide arbitrarily large indices with $v_p(P_j(n)) = 1$. An independent clock of positive rate at this prime changes that trace value when it rings: the other prime factors of this fixed value are finite in number and almost surely do not ring simultaneously. If every Galois element fixes some root in the union, only finitely many clocks can remain active while preserving every trace at all times.

Starting from $f \equiv 1$, finitely many negative prime signs do not give a normal function. Its mean is the positive product

$$\prod_{p:g(p)=-1} \frac{p-1}{p+1},$$

a finite case of (5.6). This obstruction concerns independent prime clocks acting separately and pointwise protection at all times. It does not prove that every correlated completion is impossible when the reservoir is empty. For repeated polynomial factors, even exponents should first be discarded, since they impose no sign constraint of this kind.

9 Choosing row and column weights after the prime signs

Suppose a product $m = rs$ is sampled with von Mangoldt weights on both factors. An observer who knows the random function may choose which rows and columns to emphasise. A weight depending on r times a weight depending on s is controlled by a matrix norm. A weight depending jointly on (r, s) could copy the whole sign of the summand. This distinction persists when all but one Liouville factor are kept at their original values.

Fix $0 < \alpha < \beta < 1/2$, write $L = \log X$, $u = \log L$, and let

$$\mathcal{P}_X = \{(r, s) \in \mathbb{N}^2 : X^\alpha < r \le X^\beta,\ X < rs \le 2X\}.$$

Here Λ denotes the von Mangoldt function, including its values on proper prime powers. For a complex table H define

$$\mathcal{T}_X(H) = \sup_{|a_r|,|b_s|\le 1} \left| \frac{1}{XL} \sum_{(r,s)\in\mathcal{P}_X} a_r b_s \Lambda(r)\Lambda(s) H(r, s) \right|. \tag{9.1}$$

For the constant table, $\mathcal{T}_X(1) = \beta - \alpha + o(1)$. Indeed the sum of $\Lambda(s)$ in $(X/r, 2X/r]$ is $(1 + o(1))X/r$ uniformly in the row, and partial summation gives $\sum \Lambda(r)/r = (\beta - \alpha + o(1))L$ over the row range.

Write p_j for the jth prime and put

$$\mathcal{D} = \{p_{\lfloor j\sqrt{\log\log(j+20)}\rfloor} : j \ge 3\}. \tag{9.2}$$

Let ϵ be completely multiplicative, with independent fair signs on $\mathcal{D}$ and value 1 on the other primes. All scales read the same signs. Inverting the increasing rank sequence and then using the prime number theorem and partial summation gives

$$\frac{\#\{p \le x : p \in \mathcal{D}\}}{\pi(x)} \sim (\log\log x)^{-1/2}, \qquad \sum_{\substack{p\le x\\ p\in\mathcal{D}}} \frac{1}{p} \sim 2\sqrt{\log\log x}.$$

In the second formula the integral is $\int dt/(t \log t \sqrt{\log\log t})$.

Theorem 9.1 (Separated tests on prime products)**.** *Let $G_j(m) = c_j m + d_j$, $1 \le j \le t$, be fixed pairwise non-proportional primitive integral forms with $c_j > 0$ and $d_j \ne 0$. Fix a joint table $|A(r, s)| \le 1$ before drawing ϵ. For each $0 < \gamma < 2t$, almost surely at every sufficiently large real X,*

$$\mathcal{T}_X\left(A(r, s)\prod_{j=1}^{t} \epsilon(G_j(rs))\right) \ll \exp(-\gamma\sqrt{\log\log X}). \tag{9.3}$$

The supremum in (9.1) *is taken after the draw. A single event suffices for any prescribed countable collection of these data. For tables prescribed separately at the points of the polynomial-size mesh of Appendix A, the bound holds at those points, in particular at the dyadic scales.*

For example, write $\lambda(n) = (-1)^{\Omega(n)}$ for Liouville's function and $f = \lambda\epsilon$. Take $F_1(m) = m-1$, $F_2(m) = 3m-1$, $F_3(m) = 3m-2$. With $t = 1$, $G_1 = F_1$ and $A = \prod_{i=1}^{3} \lambda(F_i(rs))$, the summand contains $f(F_1(rs))\lambda(F_2(rs))\lambda(F_3(rs))$. The last two signs are the original Liouville values. The same argument permits either other choice of the modified site.

For $t = 1$, the indicator of disagreement of $f(G_1)$ with $\lambda(G_1)$ is $(1 - \epsilon(G_1))/2$. Its mass in the normalisation of (9.1) tends to $(\beta - \alpha)/2$ by the theorem with constant coefficients and $A = 1$. Thus a change of relative prime density zero changes the observed sign on half the prime-pair mass. If a joint table could be chosen after the draw, taking $A = \prod_j \epsilon(G_j(rs))$ would give $\mathcal{T}_X(1)$ on every realisation.

Where the optimisation disappears. On a row block $R < r \le 2R$, put $V = X/R$ and give the prime columns in $(V/2, 2V]$ probability proportional to $\log s$. Cauchy–Schwarz in the columns removes b_s. In the resulting square the two rows r, r' read $2t$ affine forms at the same prime s. For $r \neq r'$ their roots are distinct away from primes dividing a fixed list of coefficients and determinants. Averaging over the independent prime signs forces even total parity at each retained prime. This projection applies to every separated test at once. Appendix A makes the local approximation uniform in the rows and takes a growing moment to reach all scales. Ford's *Poisson approximation of prime divisors of shifted primes* gives a total-variation model for one shifted prime [10]. Here the common prime column, the moving row coefficients and the optimisation in (9.1) require the joint calculation.

10 One random sign for an entire band of primes

Independent signs at individual primes provide many random choices. How many choices suffice for the same separated tests? Let $h > 0$. Band $j \ge 1$ contains primes with $jh \le \log\log p < (j+1)h$. All permitted primes in a band receive the same fair sign η_j, independently from band to band. The completely multiplicative function ξ takes value 1 at every other prime. Fix a primitive affine form $F(m) = cm + d$ with $c > 0$, $d \neq 0$. For a deterministic family $|A_X(r,s)| \le 1$ set

$$T_X = \mathcal{T}_X(A_X(r,s)\xi(F(rs))).$$

We use either all primes in each band or the primes $p \equiv 1 \bmod q_j$, where

$$q_j = 2^{\lfloor \log_2((j+64)/(\log(j+64))^2) \rfloor}. \tag{10.1}$$

A third choice, which protects a polynomial trace, is given below. All logarithms without a base are natural.

Theorem 10.1 (Band width and exceptional sign sequences)**.** *For each $h > 0$, each prescribed family of tables and each fixed $\rho > 0$,*

$$\limsup_{X\to\infty} \frac{\log \mathbb{P}(T_X > \rho)}{\log\log X} \le -\frac{\log 2}{h}. \tag{10.2}$$

Use $\log 0 = -\infty$. If $h < \log 2$, almost surely $T_{2^n} \to 0$. For a single table $A_X = A$ this convergence holds at all real scales. A countable collection fixed in advance can be treated simultaneously.

Give the space of sign sequences the metric 2^{-j} when the first differing coordinate is j. For each prescribed family, the set where T_{2^n} fails to tend to zero has Hausdorff dimension at most $\min(1, h/\log 2)$. This bound is attained within the class of dyadically varying deterministic families. If $h \ge \log 2$, one such family satisfies $\limsup_n T_{2^n} \ge \beta - \alpha$ for every *sign sequence.*

The summability threshold has a concrete origin. At $X = 2^n$ the summands use only $(\log n)/h + O(1)$ bits. The reciprocal of the number of possible prefixes is $n^{-(\log 2)/h+o(1)}$. For $A = 1$, the all-positive prefix preserves the whole prime-pair mass, so the exponent in (10.2) is already attained for this table when $0 < \rho < \beta - \alpha$. This is a lower bound on a single-scale probability. The failure at every sign sequence above the threshold uses a family designed to list prefixes on different dyades. The placement of a particular sign sequence in the exceptional set for a prescribed Liouville correlation remains open. Appendix B proves both the upper estimate and the sharp dimension statement.

For (10.1), the Siegel–Walfisz theorem [13, Corollary 5.29] and partial summation give permitted harmonic mass $(h + o(1))/\varphi(q_j)$ per band. Consequently the cumulative mass up to x has order $(\log\log\log x)^3$. The tail of the permitted set lies in 1 mod q_J after band J, so its relative prime density is zero. Its band masses have summable squares and divergent sum. The centred series of independent signed band masses therefore converges almost surely, and the flipped subcollection has asymptotically half the cumulative mass.

Corollary 10.2 (A protected quadratic trace with separated tests)**.** *In band j permit only $p \equiv -1 \bmod \max(4, q_j)$ and omit any fixed initial set of primes $p \le M_0$. Use the same $\Lambda(r)\Lambda(s)/(X \log X)$ sampling, preassigned joint tables and post-draw separated weights as in Theorem 10.1. Its probability estimate and sharp dimension statement still hold. In particular, for $h < \log 2$ the optimised tests tend to zero almost surely on the dyades for each prescribed family, and at all real scales for a fixed table. For every prescribed $f \in \mathcal{M}_{\pm}$, at every realisation,*

$$(f\xi)(n^2+1) = f(n^2+1) \quad (n \in \mathbb{Z}), \qquad (f\xi)(m) = f(m) \quad (1 \le m \le M_0).$$

Proof. The residue class is reduced, has the same eventual modulus as in (10.1), and has the same harmonic mass estimate. Appendix B uses only that estimate for the permitted primes. The lower construction uses only the number of bits. Every permitted prime is 3 mod 4, so -1 is a quadratic non-residue there. Multiplicativity gives both identities. □

In particular, one can keep every $\lambda(n^2 + 1)$ and still optimise the two coefficients of a prime-pair test on one modified Liouville site when $h < \log 2$. Theorem 1.1 treats general irreducible traces with a reservoir chosen for the polynomial. The support in the corollary was chosen for this specific trace.

11 What one cofactor bit reveals, and which twists preserve cancellation

The post-draw observation now concerns one integer. An observer knows two of its divisor labels and learns one sign from the part of a cofactor supported on small primes. The resulting absolute sums recover a specified fraction of the original bias. All the randomness in this section comes from a new sequence of independent, biased prime signs. Its law is specified below.

Fix endpoint constants $\boldsymbol{c} = (c_1, C_1, c_2, C_2)$ with $0 < c_i < C_i$. Put

$$A = X^{1/4}, \qquad B = X^{3/5}, \qquad M = X/A, \quad K = X/B, \quad U = X/(AB).$$

The labels $a \in (c_1A, C_1A]$, $b \in (c_2B, C_2B]$ are odd. Their cell is

$$I_{ab} = \{X < n \le 2X : n = am = bk - 2 \text{ for some } m, k \in \mathbb{N}\}, \qquad N_{ab} = |I_{ab}|.$$

A common divisor of a, b must divide 2. For a coprime pair this is one progression modulo ab, of length $X/(ab) + O(1) \asymp_{\boldsymbol{c}} U$. Here $U = X^{3/20}$. The sum of the cell sizes is $O_{\boldsymbol{c}}(X)$. All sums

count an origin once for each of its admissible pairs of labels. Denote a cell average by $\mathbb{E}_{ab}$, omitting empty cells.

Independently at each prime choose $R(p) = -1$ with probability

$$\delta_p = \frac{1}{4\sqrt{\log\log(e^e p)}},$$

and $R(p) = 1$ otherwise. Extend R completely multiplicatively and put $f = \lambda R$. For a cutoff $2 \le y \le X$ let $S_y(m) = \prod_{p \le y} p^{v_p(m)}$. The bit read by the observer is $B_y(n; a) = R(S_y(n/a))$. Let $g_{ab}(n)$ be a fixed complex array of modulus at most one, independent of the prime draw and of X. The original norm and the norm after reporting the bit are

$$D(g; X) = \sum_{a,b} \left| \sum_{n \in I_{ab}} \lambda(n) g_{ab}(n) \right|,$$

$$D_{\mathrm{bit}}(g; X, y) = \sum_{a,b} \sum_{e=\pm 1} \left| \sum_{\substack{n \in I_{ab} \\ B_y(n;a)=e}} f(n) g_{ab}(n) \right|. \tag{11.1}$$

The absolute values are taken separately for the two reported signs.

Theorem 11.1 (Transmission of a fixed bias)**.** *Write $u = \log\log X$ and $v_X(y) = \log(\log X/\log y)$. For almost every single draw of R, at every sufficiently large real X,*

$$\sup_{2 \le y \le X} \left| D_{\mathrm{bit}}(g; X, y) - \exp\!\left(-\frac{v_X(y)}{2\sqrt{u}}\right) D(g; X) \right| \ll_{\boldsymbol{c}} X u^{-1/4}. \tag{11.2}$$

Moreover, for $y_0 = \exp(\sqrt{\log X})$,

$$\sup_{2 \le y \le y_0} D_{\mathrm{bit}}(g; X, y) \ll_{\boldsymbol{c}} X e^{-\sqrt{u}/20}. \tag{11.3}$$

This also bounds the sum of absolute cell sums when the bit is hidden, by the triangle inequality. The starting scale may depend on the draw and the array. A single event works for any countable collection specified in advance. In particular, for $0 \le \theta \le C$ fixed in a compact interval, $y_\theta = \exp((\log X) e^{-\theta\sqrt{u}})$ gives the multiplier $e^{-\theta/2}$, uniformly in θ.

At $\theta = 1$, a cutoff $X^{o(1)}$ retains the fraction $1/\sqrt{e}$ of $D(g; X)$ up to $O(X u^{-1/4})$, even when $D(g; X)/X$ has no limit. Take $g_{ab}(n) = \lambda(n + 2)$. For any fixed finite θ, an $o(X)$ estimate after reading the bit is equivalent to the same estimate for the original two-site cell sums. The choice $g_{ab}(n) = \lambda(n + 1)\lambda(3n + 1)\lambda(3n + 2)$ gives the corresponding quartic statement. If $v_X(y)/\sqrt{u} \to \infty$, the observed norm itself is $o(X)$ for every such array. The multiplier then tends to zero. In (6.3), a character on two active irreducible factors has multiplier $e^{-4\tau}$ at local clock time τ. Formula (11.2) measures the parity left above a size cutoff in a single cofactor, under the biased sign law just defined.

11.1 Reading the bit inside one cell

Write x_+, x_- for the inner sums in (11.1). For complex numbers,

$$|x_+ - x_-| \le |x_+| + |x_-| \le |x_+ - x_-| + |x_+ + x_-|. \tag{11.4}$$

The upper bound follows by writing $2x_\pm$ as a sum or difference of $x_+ + x_-$ and $x_+ - x_-$. Set $z(n) = \lambda(n) g_{ab}(n)$ within the cell and $R_{>y}(m) = R(m)/R(S_y(m))$. Then

$$R(a)(x_+ - x_-) = \sum_{I_{ab}} z(n) R_{>y}(m), \qquad R(a)(x_+ + x_-) = \sum_{I_{ab}} z(n) R(m). \tag{11.5}$$

For example, if $m = pq$ and only p is below the cutoff, multiplying the summand by the reported sign cancels $R(p)$ and leaves $R(q)$, whose expectation is $1 - 2\delta_q$. Even powers of a prime cancel with no averaging. Thus the mean high-prime multiplier of each summand is

$$w_y(m) = \prod_{\substack{p>y \\ v_p(m) \text{ odd}}} (1 - 2\delta_p).$$

Lemma 11.2 (Damping of the unread primes)**.** *Set $y_0 = \exp(\sqrt{\log X})$ and $u = \log\log X$. Uniformly over nonempty cells and $y_0 \le y \le X$, with $v = v_X(y)$,*

$$\mathbb{E}_{ab}\left|w_y(m) - e^{-v/(2\sqrt{u})}\right| \ll_{\boldsymbol{c}} u^{-1/4}, \qquad \mathbb{E}_{ab} w_{y_0}(m) \ll_{\boldsymbol{c}} e^{-\sqrt{u}/20}.$$

Proof. Put $Z = X^{1/25}$. Count first the distinct divisors $y < p \le Z$ of m, writing H for their number, zero for an empty interval. As m runs through a progression of step b, a prime $p \nmid b$ divides m with frequency $1/p + O(1/N_{ab})$, and two such distinct primes with frequency $1/(pq) + O(1/N_{ab})$. For $p \mid b$ the frequency is zero since $am \equiv -2 \bmod b$. The removed reciprocal mass above y_0 is $O((\log X)/y_0)$. Mertens' formula and the second-moment expansion give

$$\mathbb{E}_{ab} H = v + O(1), \qquad \operatorname{Var}_{ab} H \ll 1 + v + Z^2/N_{ab} \ll_{\boldsymbol{c}} 1 + v.$$

If $y > Z$, then $v \le \log 25$, which accounts for the empty interval. A square $p^2 \mid m$ for $y < p \le Z$ occurs on a set of relative size $O(y_0^{-1} + Z/N_{ab})$. Primes above Z contribute only $O_{\boldsymbol{c}}(1)$ factors to an integer $m \ll_{\boldsymbol{c}} X$. Outside that exceptional set the number V_y of odd valuations above y differs from H by $O_{\boldsymbol{c}}(1)$.

For the primes involved, $0 \le v \le u/2$ and

$$t_p = -\log(1 - 2\delta_p) = \frac{1}{2\sqrt{u}} + O\!\left(\frac{1}{u} + \frac{1+v}{u^{3/2}}\right), \qquad t_p \gg u^{-1/2}.$$

Indeed $\log\log(e^e p)$ is between $u - v + o(1)$ and $u + O(1)$. Put $s = -\log w_y$. On the complement of the exceptional set the mean of $|s - v/(2\sqrt{u})|$, with the indicator of that complement, is

$$\ll \frac{\sqrt{1+v}}{\sqrt{u}} + \frac{1+v}{u} + \frac{(1+v)^2}{u^{3/2}}. \tag{11.6}$$

For $v \le \sqrt{u}$ this suffices since e^{-s} is 1-Lipschitz on $s \ge 0$. For $v > \sqrt{u}$, discard also $H < v/2$, of relative size $O(1/v)$ by Chebyshev. On the remainder both s and $v/(2\sqrt{u})$ are $\gg v/\sqrt{u}$. The mean-value theorem supplies the factor $e^{-cv/\sqrt{u}}$ in front of (11.6). Writing $v = t\sqrt{u}$ shows that the product is $O(u^{-1/4})$ uniformly for $t \ge 1$. Each discarded set costs at most its relative size in the comparison of two numbers in $[0, 1]$.

For the second estimate keep only $y_0 < p \le Z_0 = \exp((\log X)^{3/4})$ with $p \nmid b$. A square at any of these primes costs $O(y_0^{-1} + Z_0/N_{ab})$. Elsewhere $w_{y_0}(m)$ is at most $\prod_p (1 - 2\delta_p \mathbf{1}_{p|m})$ over this window. Its cell average is approximated by $\prod_p (1 - 2\delta_p/p)$. Indeed, truncate inclusion–exclusion at the two consecutive orders $\ell - 1, \ell$, with $\ell = 2\lceil u \rceil$. The Bonferroni inequalities apply also to these factors in $[0, 1]$, by independently thinning each incidence with probability $2\delta_p$. The CRT error of each truncated expectation is $O((2Z_0)^\ell/N_{ab}) = X^{-3/20+o(1)}$. The two model bounds differ by at most $J^\ell/\ell!$, where

$$J = \sum_{\substack{y_0 < p \le Z_0 \\ p \nmid b}} 2\delta_p/p \;\ge\; (\tfrac{1}{8} + o(1))\sqrt{u}, \qquad J \ll \sqrt{u}.$$

Here the harmonic mass of the window is $u/4 + O(1)$ and the contribution of $p \mid b$ is negligible. The model product is $\le e^{-J}$ and the truncation error is smaller than $e^{-\sqrt{u}}$. This gives the stated exponential bound. □

Proof of Theorem 11.1. For $w \geq y_0$ consider a cell average of $z'(n)R_{>w}(m)$, where $|z'| \leq 1$ is deterministic after conditioning on the signs at $p \leq w$. Its conditional mean is $\mathbb{E}_{ab} z' w_w(m)$. Changing $R(p)$ for $p > w$ affects the average by at most $2\gamma_p$, where $\gamma_p = \mathbb{P}_{ab}(p \mid m) \leq p^{-1} + N_{ab}^{-1}$. The preceding count and the bound on factors above Z give $\sum_{p>w} \gamma_p \ll 1 + u$. Thus

$$\sum_{p>w} \gamma_p^2 \ll_{\boldsymbol{c}} (1+u)(y_0^{-1} + N_{ab}^{-1}).$$

The classical bounded-differences inequality [16], separately for real and imaginary parts, bounds the probability of an error exceeding $\sigma = e^{-\sqrt{u}/20}$ by

$$4\exp\left(-\frac{c_{\boldsymbol{c}}\sigma^2}{(1+u)(y_0^{-1} + N_{ab}^{-1})}\right).$$

This bound is independent of the conditioned signs. On $T \leq X \leq 2T$ there are $T^{O(1)}$ possibilities for the two integer labels, the interval endpoints and the last primes up to any two cutoffs. The array $g_{ab}(n)$ stays fixed on each such choice. It suffices to use $e^{-\sqrt{\log\log(2T)}/20}/2$ as the error threshold and the lower bounds for y_0, N_{ab} on the block. The union bounds are summable over dyadic T since $\exp(\sqrt{\log T})$ exceeds every power of $\log T$. Borel–Cantelli gives all the following comparisons on one event. Only the three specified choices of z' below enter this argument.

For $y \geq y_0$, apply the comparison at $w = y$ with $z' = z$. Lemma 11.2 approximates the first normalised sum of (11.5) by $e^{-v/(2\sqrt{u})}\mathbb{E}_{ab} z$ to error $O_{\boldsymbol{c}}(u^{-1/4})$. For $y < y_0$ condition instead at $w = y_0$ and include the factors of R on $y < p \leq y_0$ in z'. Its conditional mean has modulus at most $\mathbb{E}_{ab} w_{y_0} \ll_{\boldsymbol{c}} \sigma$. Here the proposed multiplier is $\leq e^{-\sqrt{u}/4}$. Finally use $w = y_0$ with $z' = zR(S_{y_0}(m))$ for the second sum of (11.5). Its normalised modulus is also $O_{\boldsymbol{c}}(\sigma)$. Applying (11.4) and summing with weights N_{ab} proves (11.2). For $y \leq y_0$ both absolute values on the right of (11.4) are $O_{\boldsymbol{c}}(N_{ab}\sigma)$ by the comparison at $w = y_0$, which gives (11.3). □

11.2 When a multiplicative perturbation transfers cancellation

The prime draw in Theorem 11.1 changes a proportion tending to zero of the primes. Nevertheless, almost surely

$$\sum_{\substack{p \leq x \\ R(p) = -1}} \frac{1}{p} = \left(\tfrac{1}{2} + o(1)\right)\sqrt{\log\log x}. \tag{11.7}$$

Indeed the centred reciprocal series has summable variances, and Mertens' formula with partial summation evaluates $\sum_{p \leq x} \delta_p/p$. The centred indicators along the sequence of primes obey the strong law, and $\delta_p \to 0$, which gives the relative-density assertion. The next theorem identifies a class of changes for which cancellation of the original product *can* be transferred in both directions.

For functions F, G of modulus at most one define

$$D_{F,G}(X; \boldsymbol{c}) = \sum_{a,b} \left| \sum_{n \in I_{ab}} F(n)G(n+2) \right|.$$

The moduli are taken before summing over the two odd divisor labels. A dilation by (d, e) multiplies both endpoints of the a-interval by d and both endpoints of the b-interval by e. Each dilation in the next statement is fixed as $X \to \infty$.

Theorem 11.3 (Stability on the family of dilated boxes)**.** *Let u_1, u_2 be functions of modulus one that are multiplicative on coprime arguments. Assume for $i = 1, 2$ that*

$$u_i(2^j) = 1 \quad (j \geq 1), \qquad E_i(Y) = \sum_{p > Y} \frac{|1 - u_i(p)|^2}{p} \longrightarrow 0. \tag{11.8}$$

The values at higher odd prime powers are arbitrary. For fixed bounded $|F|, |G| \le 1$ one has

$$D_{Fu_1,Gu_2}(X;\boldsymbol{c}) = o_{\boldsymbol{c}}(X) \quad \text{for every fixed } \boldsymbol{c} \\ \iff D_{F,G}(X;\boldsymbol{c}) = o_{\boldsymbol{c}}(X) \quad \text{for every fixed } \boldsymbol{c}.$$

In either direction, to obtain the assertion on a specified box it suffices to assume the other estimates on all of its fixed positive odd integral dilations (d, e), including $(1, 1)$.

Taking $F = G = \lambda$ gives the comparison for ordinary multiplicative $f_i = \lambda u_i$. The sum in (11.8) is twice the usual squared pretentious distance to 1 over the indicated primes. For signs, it is four times the sum of $1/p$ over flipped primes. Equation (11.7) explains why cancellation of an unread perturbed product cannot be passed back to Liouville by this theorem. Complex phases may satisfy (11.8) even when their angles have a divergent first weighted moment.

Proposition 11.4 (A quantitative estimate on one box)**.** *For a fixed box choose C so that $m \le CM$ and $k \le CK$. Let $v_{1,X}, v_{2,X}$ be completely multiplicative functions of modulus one, possibly varying with X, and put $L_1 = M$, $L_2 = K$ and*

$$\varepsilon_i^2 = \sum_{p \le CL_i} \frac{|1 - v_{i,X}(p)|^2}{p}.$$

For $|F|, |G| \le 1$ one has, uniformly in the functions,

$$\frac{|D_{Fv_{1,X},Gv_{2,X}}(X;\boldsymbol{c}) - D_{F,G}(X;\boldsymbol{c})|}{X} \ll_{\boldsymbol{c}} \varepsilon_1 + \varepsilon_2 + o(1).$$

All prime values, including the ones at 2, enter this estimate.

Small *total* distance up to the cofactor scales gives continuity on one box. Finite distance with an arbitrarily large initial contribution uses a different argument. A fixed small odd prime may appear in a cofactor. Testing $d \mid m$ changes the label from a to ad, so the input to Theorem 11.3 naturally includes other boxes. Appendix C proves both assertions.

11.3 The prime two and genuine Möbius cofactors

For functions s, t of the cofactors let

$$C_{s,t}(X;\boldsymbol{c}) = \sum_{a,b} \left| \sum_{n=am=bk-2 \in I_{ab}} s(m)t(k) \right|.$$

A superscript odd restricts origins to odd n. Complete multiplicativity gives $D_{\lambda,\lambda} = C_{\lambda,\lambda}$, also with this restriction. For even n, one of $n, n+2$ is divisible by four. Because both labels are odd, the same holds for one of the cofactors and $\mu(m)\mu(k) = 0$ exactly. Each occupied cell has odd step ab, so the odd layer carries half the total incidence mass up to $O_{\boldsymbol{c}}(AB)$.

Proposition 11.5 (Transfer to squarefree cofactors)**.** *The two families of boxes in Theorem 11.3 satisfy the equivalence*

$$[\ \forall \boldsymbol{c},\ D^{\mathrm{odd}}_{\lambda,\lambda}(X;\boldsymbol{c}) = o_{\boldsymbol{c}}(X)\ ] \quad \iff \quad [\ \forall \boldsymbol{c},\ C_{\mu,\mu}(X;\boldsymbol{c}) = o_{\boldsymbol{c}}(X)\ ].$$

For both parities of the origin the equivalence holds with $D_{\lambda,\lambda}$ and $C_{\mu^{[2]},\mu^{[2]}}$, where $\mu^{[2]}(2^j w) = (-1)^j \mu(w)$ for w odd.

The square-divisor expansions proving this result appear at the end of Appendix C. Thus an estimate for ordinary Möbius cofactors on odd divisor labels has an exact target on the Liouville side. Questions about recovering the labels or the displacement from partial observations are treated in *Observations and Shift Decoding in Divisor Incidences* [21].

12 Periodic observations and fluctuations on ordinary intervals

A change at each of many primes can leave a measurable fraction of the original bias. We first ask how finely the origin may be identified while this fraction remains visible. We then compare two experiments that record fluctuations about the noise mean.

12.1 How large can the observed period be?

Let X be a positive integer, $I = \{X+1, \dots, 2X\}$ with uniform law, and fix $0 < a < b < 1$. Define

$$R = \exp((\log X)^a), \qquad Z = \exp((\log X)^b), \qquad \mathcal{B} = \{p : R < p \le Z\}, \qquad H = \sum_{p \in \mathcal{B}} p^{-1}.$$

Here and below p is prime. Mertens' theorem gives $H \sim (b-a)\log\log X$. Fix $t > 0$. At this scale draw independent signs ξ_p with $\mathbb{P}(\xi_p = -1) = t/H$ on $\mathcal{B}$, for X sufficiently large, and take $\xi_p = 1$ elsewhere. Extend ξ completely multiplicatively. The law is specified separately at each scale. In Section 10 all permitted primes of a band share one sign.

Let P be a product of $d \ge 1$ fixed primitive affine forms with positive slopes and distinct rational roots. For any deterministic complex array $|v_X(n)| \le 1$ write

$$C_q(r) = X^{-1} \sum_{\substack{n \in I \\ n \equiv r \bmod q}} v_X(n), \qquad C_q^{\xi}(r) = X^{-1} \sum_{\substack{n \in I \\ n \equiv r \bmod q}} v_X(n)\xi(P(n)).$$

All forms are positive on I for large X. Each residue carries its share of the whole interval. For $v_X = \lambda(P)$, C_q^{ξ} is the signed mass of the perturbed source $\lambda\xi$ on that class.

Theorem 12.1 (A sharp range for periodic tests)**.** *Put $\rho_d = e^{-2dt}$. For each prescribed array and every $D > 0$, with probability $1 - O_D(X^{-D})$, simultaneously for $1 \le q \le X$,*

$$\sum_{r \bmod q} |C_q^{\xi}(r) - \rho_d C_q(r)| \ll_{P,a,b,t} H^{-1/2} + \sqrt{q/X}. \tag{12.1}$$

Consequently the error tends to zero uniformly for $q \le Q_X = o(X)$. Here Q_X can be any integer bound. The period and the complex coefficients of modulus at most one on its residues may be chosen after the draw. The condition $Q_X = o(X)$ is necessary for convergence for every prescribed array, when $1 \le Q_X \le X$.

For $P(n) = n(n+1)(3n+1)(3n+2)$, the multiplier is e^{-8t}. The exponent is the classical degree-four Walsh multiplier with one-site correlation e^{-2t} [19, Proposition 8.28]. The content of (12.1) is its comparison with the empirical signed masses uniformly over periods. For example $Q_X = \lfloor X/\log\log X \rfloor$ gives error $O((\log\log X)^{-1/2})$. Duality here concerns a single preassigned array and all bounded functions of the observed residue.

The range $q = o(X)$ also occurs in *Multiplicative functions in short arithmetic progressions.* There Klurman, Mangerel and Teräväinen bound the variance over reduced residues of a bounded multiplicative function, with its character main term, for almost all moduli [15, Theorem 1.4]. In (12.1) the comparison is with the same array on the same residues before the prime flips. Concentration over the flips permits every modulus and all its residue classes for this comparison.

Proof. Let $T(n) = \{p \in \mathcal{B} : v_p(P(n)) \text{ odd}\}$. All sufficiently large p have d distinct simple roots. Counting in I gives $d/p + O_P(1/X)$ for divisibility by p, and $d^2/(ps) + O_P(1/X)$ at distinct p, s. Hence the count of prime divisors in the band has squared deviation from dH of mean $O_P(H + Z^2/X)$. Replacing divisibility by odd valuation changes the count by at most the

number of $p^2 \mid P(n)$ in the band. The second moment of this number is $O_P(1)$, by the same count modulo p^2 and $p^2 s^2$. Thus

$$\mathbb{E}_I(|T(n)| - dH)^2 \ll_P H, \qquad \mathbb{E}_I\big|(1-2t/H)^{|T(n)|} - \rho_d\big|^2 \ll_{P,t} H^{-1}. \tag{12.2}$$

For the second bound use $-\log(1-2t/H) = 2t/H + O_t(H^{-2})$ and the Lipschitz bound for the exponential on the negative half-line. The noise mean of $\xi(P(n))$ is the product on the left. Summing the errors of the class means therefore costs $O(H^{-1/2})$ in total, uniformly in q.

It remains to centre the random class vector about its noise mean. For a fixed q, denote by $\widetilde{C}_q$ the vector obtained by setting the signs of primes $p \mid q$ to $+1$. This changes the vector, in the sum of moduli over residues, by at most $O_P(\log X/(R\log R))$, even when every one of these signs is reversed. Indeed q has at most $\log X/\log R$ prime divisors in the band, and at most $O_P(X/p)$ origins involve any one of them. The same bound holds for the mean vector.

For $p \nmid q$ let $m_{p,r} = X^{-1}\#\{n \in I : n \equiv r \bmod q,\ p \in T(n)\}$. Counting CRT classes and then forgetting the residue give

$$\max_r m_{p,r} \ll_P (pq)^{-1} + X^{-1}, \qquad \sum_r m_{p,r} \ll_P p^{-1}, \qquad q\sum_r m_{p,r}^2 \ll_P p^{-2} + \frac{q}{Xp}. \tag{12.3}$$

Resampling ξ_p changes the coordinate at r by at most $2m_{p,r}$, with probability at most $2t/H$. The Efron–Stein inequality in its independent-coordinate form [6, 24], followed by Cauchy–Schwarz over the q coordinates, gives

$$\mathbb{E}_\xi \sum_r |\widetilde{C}_q(r) - \mathbb{E}_\xi \widetilde{C}_q(r)| \ll_{P,t} \left(\frac{1}{H}\sum_{p\in\mathcal{B}}\left(p^{-2} + \frac{q}{Xp}\right)\right)^{1/2} \ll_{P,t} R^{-1/2} + \sqrt{q/X}.$$

The bound for complex masses follows by applying the variance inequality to their real and imaginary parts. The $q/(Xp)$ term in (12.3) permits periods larger than X/Z, where counting a full period of length pq inside I ceases to be possible.

For each fixed q, changing one sign in the original sum of moduli on the left of (12.1) changes it by $O_P(1/p)$. McDiarmid's inequality [16] gives probability at most $\exp(-c_P R/H)$ of an excess $H^{-1/2}$ above its expectation. A union bound over the X moduli proves the estimate. The exceptional probabilities are summable on the integer scales for any coupling of their marginal draws.

Finally fix any integer $q \le X$ and set $v_X(n) = \mathbf{1}_{X<n\le X+q}$. Each residue contains exactly one selected origin. On every draw the left side of (12.1) is at least $(1-\rho_d)q/X$. Choosing $q = Q_X$ on a subsequence with Q_X/X bounded below proves necessity for this class of arrays. □

12.2 Words with a growing number of sites

Let $k = k(X) \ge 1$. At each origin prescribe a word $w(n) \in \{-1,1\}^k$ before the draw. Observe

$$w^\xi(n) = \big(w_i(n)\xi(n+i)\big)_{0\le i<k}.$$

A source word $w_i(n) = f(n+i)$ is one instance. Denote by μ_q^ξ the empirical probability of $(n \bmod q, w^\xi(n))$ for $n \in I$. For its comparison law ν_q, sample n uniformly, keep its residue, and flip each coordinate of $w(n)$ independently with probability $\pi_t = (1-e^{-2t})/2$.

Proposition 12.2. *If $k = o(H)$ and $2^k k^2 Q_X = o(X)$, then $\sup_{q\le Q_X} d_{\mathrm{TV}}(\mu_q^\xi, \nu_q) \to 0$ in probability. In particular every range $q \le X^{1-\delta}$ with fixed $\delta > 0$ is allowed. All tests of the pair consisting of the residue and the word may be chosen after the draw.*

Proof. For a fixed origin the k sites involve disjoint prime sets because $R > k$. Their noise signs are independent given the origin. The mean squared error of each individual flip probability from π_t is $O_t(1/H)$, by the one-form proof of (12.2), uniformly for $0 \le i < k$: every shifted interval has the same length and one root modulo each prime power. For products of Bernoulli laws of parameters s_i and π_t, the squared total variation distance is at most $C_t \sum_i (s_i - \pi_t)^2$. Indeed squared Hellinger distances in the convention $\sum(\sqrt{P} - \sqrt{Q})^2$ are subadditive for products by multiplication of affinities, dominate squared total variation, and have this quadratic bound against a fixed interior parameter. Averaging over the origins and retaining the residue gives $d_{\mathrm{TV}}(\mathbb{E}_\xi \mu_q^\xi, \nu_q) \ll_t \sqrt{k/H}$.

In (12.3) use the union of the k sites. The first two bounds acquire a factor k and the last a factor k^2. For a given residue, changing a sign redistributes at most that vulnerable mass among the 2^k states. The sum of the squared changes in their probabilities is at most $4m_{p,r}^2$, with $m_{p,r}$ now counting the union. Sum the coordinate variances first and apply Cauchy–Schwarz over $q2^k$ entries. The expected centred variation is

$$\ll_t 2^{k/2} k(R^{-1/2} + \sqrt{q/X}) + k \log X/(R \log R).$$

The last term again removes the primes dividing q. One prime can change total variation by $O(k/p)$. The union bound for an excess $H^{-1/2}$ is at most $X \exp(-cR/(k^2 H))$. Under the stated hypotheses all the errors tend to zero. The target retains the joint distribution of the source word, including its dependence on the residue. □

12.3 Two experiments for the same signed mean

For clarity take four forms, $d = 4$. Keep the interval I, the prime band $\mathcal{B}$ and its harmonic mass H from Section 12. Fix a real array $|v(n)| \le 1$, allowed to depend on X, before drawing the noise. Write

$$M = \mathbb{E}_I v(n), \qquad A^\xi(v) = \mathbb{E}_I v(n)\xi(P(n)), \qquad F^\xi(v) = A^\xi(v) - \mathbb{E}_\xi A^\xi(v).$$

The two experiments below use different probabilities from the choice t/H above. Both can be applied to $v(n) = \lambda(P(n))$. Their output then depends on the unknown ordinary quartic mean on I.

Two classical comparisons help interpret these statistics. Harper's low-moment estimates for ordinary sums of Steinhaus and Rademacher random multiplicative functions separate their typical and root-mean-square scales [12, Theorems 1–2]. Soundararajan and Xu give Gaussian limits for sums of a Steinhaus function over suitable sets of integers [23, Theorem 1.1]. Here the choice of flip probabilities controls the first-order contribution of each prime, with the signed mean of a prescribed array as its coefficient after averaging over the chosen interval.

Theorem 12.3 (Rare and Gaussian fluctuations)**.** (i) *Flip each prime in $\mathcal{B}$ independently with probability $p/(4Z)$. Uniformly for the prescribed arrays,*

$$\frac{Z}{16H} \operatorname{Var}_\xi A^\xi(v) = M^2 + O(H^{-1/2}), \qquad \sqrt{\frac{Z}{16H}} F^\xi(v) \longrightarrow 0 \quad \textit{in probability.} \tag{12.4}$$

The normalised covariance $\frac{Z}{16H} \operatorname{Cov}_\xi(A^\xi(v), A^\xi(1))$ equals $M + O(H^{-1/2})$.

(ii) *Retain the full dyadic intervals $I_j = (2^j, 2^{j+1}] \subset (R, Z]$. Put $H_j = \sum_{p \in I_j} 1/p$, $K_j = \sum_{p \in I_j} 1/p^2$, $H' = \sum_j H_j$. Choose j with probability H_j/H'. Given j, flip primes in I_j independently with probability $1/4$, and keep all other primes fixed at sign $+1$. With $m = 1/2$, $\sigma = \sqrt{3}/2$, set*

$$Y_j(v) = \frac{A^\xi(v) - \mathbb{E}_\xi[A^\xi(v) \mid j]}{4\sigma\sqrt{K_j}}, \qquad G_j = \frac{1}{\sigma\sqrt{K_j}} \sum_{p \in I_j} \frac{\xi_p - m}{p}.$$

Then

$$\mathbb{E}_{j,\xi}|Y_j(v) - MG_j|^2 \le \frac{\mathrm{Var}_I(v)}{2H'} + o(1/H'). \tag{12.5}$$

If $M \to m_0$ along any sequence of scales, $(Y_j(v), G_j)$ converges in law to $(m_0 G, G)$ for a standard real Gaussian G. The result includes $m_0 = 0$.

In the first experiment the normalised second moment can stay positive even though a typical draw has vanishing normalised fluctuation. In the second, the fluctuation has a reference G_j computed solely from the prime signs. A pair of draws on the same interval also removes the noise expectation, which otherwise depends on v.

Corollary 12.4 (Recovering the signed mean from two draws)**.** *Condition on the sampled j in part* (ii) *and take two independent sign vectors ξ, ξ' there. Define*

$$\widehat{M} = \frac{A^{\xi}(v) - A^{\xi'}(v)}{4\sum_{p\in I_j}(\xi_p - \xi'_p)/p}, \tag{12.6}$$

assigning zero if the denominator vanishes. For $0 < \varepsilon \le 1$,

$$\mathbb{P}_{j,\xi,\xi'}(|\widehat{M} - M| > \varepsilon) \ll (H'\varepsilon^2)^{-1/3} + \sqrt{\log(2R)/R}.$$

In particular $\widehat{M} - M \to 0$ in probability even if M has no limit. The inputs in (12.6) *are the two full perturbed averages on the same set of origins.*

Here is a finite example of the sign and factor 4 in (12.6). Take $180 < n \le 360$, $P(n) = n(n+1)(3n+1)(3n+2)$, and the five primes in $(16, 32]$. Exactly 69 origins have no active prime of odd valuation. For the five singleton sets in increasing prime order the counts are $21, 17, 17, 12, 7$. Put $M = -1/20$. On the singleton for p distribute total normalised weight $4M/p$ equally, put the remaining mass $M(1 - 4\sum_p 1/p)$ on the empty set, and give zero weight to multiple hits. All individual weights have modulus below one. Then $A^{\xi}(v) = M(1 - 4\sum_p 1/p) + 4M\sum_p \xi_p/p$, so every pair with nonzero denominator returns the negative mean exactly.

Appendix D gives the estimates. The arithmetic step is a bound on how often the mean seen at an odd prime valuation can differ from M on a large set of primes. Product Fourier analysis then identifies which part of the noise sees those conditional means. In particular (12.6) can recover a signed bias from sufficiently accurate observations while leaving its arithmetic value undetermined.

13 What passes back to the source

A trace preserved pointwise may have an unknown mean. The completion theorem transfers all of its values and leaves that arithmetic question open. Its affine estimates hold on the specified intervals at positive times or after the stated delays. The empirical law of the marks may itself move with the scale.

A post-draw test depends on what the observer can distinguish. On prime products the permitted choice factors between the two prime variables. In divisor cells the reported parity is taken inside the first cofactor, before the absolute values are summed. For $g_{ab}(n) = \lambda(n+2)$, Theorem 11.1 shows that cancellation after this reading at a finite critical parameter is equivalent to cancellation for the original pair in the same norm. When $v_X(y)/\sqrt{\log\log X} \to \infty$ the multiplier tends to zero for every initial bias. Theorem 11.3 supplies another route back to the original pair if the change has finite quadratic prime distance and the family of dilated boxes is available. Establishing the required Liouville cancellation on those boxes remains an arithmetic problem.

For the ordinary average M on one dyadic interval, (12.1) retains its periodic tests with a known attenuation. In (12.4), large but rare fluctuations can account for the second moment. In (12.6) two measurements of whole perturbed averages on a common random band recover the signed M in probability. For $v = \lambda(P)$, the size of those measured averages at the fluctuation scale is itself arithmetic information.

A Prime columns and the moments of a row pair

The input from the distribution of primes is a short-range divisibility model. We prove exactly the version needed for the two matrix arguments. All logarithmic savings below are arbitrary fixed powers. In particular the order of a subsequent matrix moment never enters the dimension of this model.

Lemma A.1 (Divisibility at a prime centre)**.** *Let V lie between two fixed positive powers of X, $0 < b < 1$, $L = \log X$, $u = \log L$, $Y \ge \exp(u^2)$ and $z = \exp(L^b)$. Draw a prime $s \in (V/2, 2V]$ with probability proportional to $\log s$. At each prime $Y < p \le z$ in any specified subset, let I_p indicate membership of $s \bmod p$ in a prescribed set of at most d reduced classes, with d fixed. Up to $O_{b,d,D}(L^{-D})$ in total variation, the I_p are independent Bernoulli variables with parameters equal to the numbers of classes divided by $p-1$, for every fixed $D > 0$. The implied constant may also depend on the two powers bounding V.*

For a fixed number of affine forms with invertible slopes at these primes, one may disregard occurrences of p^2 dividing any form at the same total error. Classes and slopes may vary with the rows and with X.

Proof. Write μ for the law of s and $E(q) = \max_{(r,q)=1} |\mu(s \equiv r \bmod q) - 1/\varphi(q)|$. The Bombieri–Vinogradov theorem [4, 13], in the form of [13, Section 17.3], gives $\sum_{q \le V^{1/3}} E(q) \ll_D L^{-D}$. One can use the two endpoints separately, remove proper prime powers at total normalised cost $O(V^{-1/6} \log^2 V)$ over these moduli, and normalise by $\sum_{V/2 < s \le 2V} \log s \asymp V$. PNT with a logarithmic saving absorbs the normalising error times $\sum_{q \le V^{1/3}} 1/\varphi(q)$. The trivial bound $E(q) \ll L/q$ and Cauchy–Schwarz imply for fixed $K_0 \ge 1$

$$\sum_{q \le V^{1/3}} K_0^{\omega(q)} E(q) \ll_{K_0,D} L^{-D}. \tag{A.1}$$

Indeed $\sum_{q \le V^{1/3}} K_0^{2\omega(q)}/q \ll L^{O_{K_0}(1)}$ by the Euler product, and the power in the unweighted theorem can be chosen afterwards.

For a subset of primes with product q, the Chinese remainder theorem gives at most $d^{\omega(q)}$ reduced residue classes. Its joint incidence probability hence differs from the independent product by at most $d^{\omega(q)} E(q)$. To apply this to an arbitrary bounded test of the vector, expand by differences from the all-zero vector, keeping subsets of size less than $\ell = \lceil u^2 \rceil$. Each coefficient has modulus $\le 2^{\omega(q)}$. All resulting products, including those of size ℓ, satisfy $q \le z^\ell = V^{o(1)}$. Formula (A.1) controls the sum of their errors.

If $W = \sum I_p$, the truncated expansion is exact at $W < \ell$. Otherwise its error is at most $(1 + 3^\ell)\binom{W}{\ell}$, by $\binom{W}{s} \le \binom{W}{\ell}\binom{\ell}{s}$ for $s \le \ell$. In the model $\mathbb{E}\binom{W}{\ell} \le (C_d u)^\ell/\ell!$ by Mertens' formula. When transferring this last factorial moment the factor 3^ℓ is absorbed as $3^{\omega(q)}$ in (A.1), since these terms use exactly ℓ primes. The model remainder is smaller than every L^{-D}. This proves the first assertion. For the second, each root lifts to exactly one class modulo p^2. Sum the model probabilities $O_d(1/\varphi(p^2))$ and the errors $O_d(E(p^2))$. Since $z^2 < V^{1/3}$ eventually, the total is $O_d(1/Y) + O_{d,D}(L^{-D})$. □

Proof of Theorem 9.1. The contribution of proper prime powers in the row or the column is $O_{\alpha,\beta}(X^{-\kappa})$ in (9.1), for some $\kappa > 0$. For either variable w, sum $O((X/w+1)L^2)$ over its proper powers, using $\#\{p^j \le W : j \ge 2\} \ll W^{1/2}\log W$. The lower bounds are X^{α} and $X^{1-\beta}$ respectively, and both variables are $\ll X^{1-\delta}$ for fixed $\delta > 0$. We henceforth use primes with their logarithmic weights.

Split the row range into $O(L)$ dyadic blocks. For each use the full interval $(R, 2R]$ with possible endpoint cuts included in the matrix entries. Put $V = X/R$, μ as in Lemma A.1, and $w_r = (\log r)/R$. Removing the column coefficient by Cauchy–Schwarz bounds the block sum divided by X by a constant times

$$\left(\sum_{r,r'} w_r w_{r'} |D_{rr'}|\right)^{1/2}, \quad D_{rr'} = \int U_{rr'}(s)\epsilon\left(\prod_{j=1}^{t} G_j(rs)G_j(r's)\right) d\mu(s). \tag{A.2}$$

Here $|U_{rr'}| \le 1$ contains the joint tables and the window cuts of both rows. This upper bound is independent of both optimised coefficients. We have $\sum w_r \ll 1$ and $\sum w_r^2 \ll (\log R)/R$.

For distinct large prime rows the $2t$ forms have distinct rational roots. A coincidence between rows would force r/r' to belong to a fixed finite list of rational numbers, which is impossible for distinct primes both tending to infinity. The product of the nonzero slopes, constants and pairwise determinants has logarithm $O_t(L)$ for the specified forms. Discard primes dividing it. In $\mathcal{D} \cap (Y, z]$, with $Y = \exp(u^2)$, the remaining harmonic mass is

$$H = (2\sqrt{b} + o(1))\sqrt{u}$$

uniformly in the rows; the deletion costs $O_t(L/Y)$ and the initial segment has mass $O(\sqrt{\log u})$.

Take $2k$ independent columns for the $2k$th absolute moment of $D_{rr'}$. Integration over the prime signs annihilates any tuple with odd total valuation at a fair prime. After this projection discard the bounded weights and impose parity only at the retained primes. In each copy Lemma A.1 models the incidence of the $d = 2t$ distinct roots with parameter $q_p = d/(p-1)$. The square exclusions allow valuation parity to be replaced by divisibility, and the product of $2k$ laws costs $O_D(kL^{-D})$ in total variation. Thus

$$\mathbb{E}|D_{rr'}|^{2k} \le \prod_{p \text{ retained}} \frac{1 + (1-2q_p)^{2k}}{2} + O_D(kL^{-D}).$$

For $k = O(u)$ the logarithm of the product is $-4ktH + O_t(k^2/Y)$. For $k = \lceil B_0\sqrt{u} \rceil$ with B_0 fixed, choose D sufficiently large. Taking the $2k$th root gives $\|D_{rr'}\|_{2k} \ll e^{-2tH}$, with H here the actual retained mass; its lower asymptotic is uniform. On the diagonal $|D_{rr}| \le 1$. Minkowski in (A.2), followed by summation over blocks and division by L, gives for the left side of (9.3)

$$\|\mathcal{T}_X\|_{4k} \ll_{B_0} e^{-(2t\sqrt{b}+o(1))\sqrt{u}} + X^{-\alpha/3} + O(X^{-\kappa}).$$

For any $\gamma' < 2t\sqrt{b}$, Markov's inequality bounds the probability of exceeding $e^{-\gamma'\sqrt{u}}$ by $O_D(L^{-D})$ for any prescribed D. Choose B_0 first using the strict gap in the exponent, then choose the logarithmic saving in the column lemma.

In $[2^m, 2^{m+1}]$ take a mesh with relative spacing at most m^{-4} and $O(m^4)$ points, including its endpoints. Borel–Cantelli yields the bound at all sufficiently large mesh points. Between neighbours, a boundary of the product window costs at most $O(L^2)\sum_r (m^{-4}X/r + 1)$ in unnormalised positive mass, summed over the row range. A boundary of the row window near $W = X^{\alpha}$ or X^{β} costs $O(L^2(m^{-4}W + 1)X/W)$. After division by XL these are $O(m^{-2})$ plus power savings. Normalisation changes by $O(m^{-4})$ times a bounded mass. This extends the estimate for an unchanged joint table to every real X. Given $\gamma < 2t$, choose $b < 1$ and γ' with $\gamma < \gamma' < 2t\sqrt{b}$. Countable intersections finish the proof. □

B Band parities, summability and Hausdorff dimension

Proof of Theorem 10.1. Use the prime-power removal and row blocks of Appendix A. Give the full prime row and column intervals probability measures proportional to their logarithmic weights. The matrix on a block is $K(r,s) = A_X(r,s)\xi(F(rs))$ with the cutoffs inserted. Its contribution to $\mathcal{T}_X$ is at most its operator norm from $L^2(s)$ to $L^2(r)$ times a deterministic weight. The block weights have bounded sum, since each is $O(1/\log X)$. For $m \geq 2$,

$$\|K\|_{\mathrm{op}}^{2m} \leq \mathrm{tr}((KK^*)^m).$$

In the trace expansion choose $r_1, \dots, r_m$ and $s_1, \dots, s_m$ independently from their respective laws, with $r_{m+1} = r_1$. Column s_i reads the two entries at (r_i, s_i) and (r_{i+1}, s_i). Averaging the trace over the band signs requires even total valuation in every band. Take absolute values of the remaining coefficients, and keep only the full bands lying in $au \leq \log\log p < bu$ with $0 < a < b < 1$. Their number is $J = (b-a)u/h + O(1)$.

Adjacent equal rows cost $O(m \log R/R)$. Fix the remaining row tuple, so $r_i \neq r_{i+1}$. Primes dividing $cd \prod_i r_i(r_i - r_{i+1})$ are exceptional. Above $\exp(L^a)$ their reciprocal sum is $O_{c,d}(mLe^{-L^a})$. With probability $1 - O_D(mL^{-D} + m^2Le^{-L^a})$ over the columns, no exceptional prime in the window divides either entry at any column. Indeed each has at most two reduced roots there (eventually all these primes exceed the fixed $|cd|$ and are below the row primes). Sum their one-prime probabilities using (A.1). This deletion matters for shared signs, since an extra hit inside a band could reverse its parity.

For each column pair the good roots are distinct. Lemma A.1 and independence of the columns transfer their m divisibility profiles jointly at cost $O_D(mL^{-D})$, including squares. In this model the character of total parity at a good permitted prime has mean $(1 - 4/(p-1))^m$. Different primes are independent. Write H_j for the harmonic mass of the permitted primes in band j. The modulus of the mean band character is bounded by

$$\exp(-4mH_j + O_{c,d}(m^2Le^{-L^a})).$$

The local probability of even parity is one half of one plus this mean. The product over J bands bounds the trace moment up to the preceding additive errors.

For the two residue-class choices, Siegel–Walfisz applies uniformly since $q_j \ll j$ and the primes here have $\log p$ exponential in j. Partial summation between the band endpoints gives $H_j = h/\varphi(q_j) + O_{h,D}(j^{-D})$ for every D (with $\max(4, q_j)$ for Corollary 10.2). On the retained bands $H_j \gg_{a,b,h} (\log u)^2/u$. With all primes $H_j = h + o(1)$. Put $m = \max(2, \lfloor u/(\log u)^{3/2} \rfloor)$. The minimum of mH_j tends to infinity and $m = o(u)$. Thus the product is at most $\exp(-J\log 2 + o(u))$, uniformly over the tables and row tuples. The transfer and collision errors are smaller upon choosing a sufficiently large fixed D. Minkowski over blocks and Markov at fixed ρ cost $\exp(O_\rho(m)) = \exp(o(u))$. The negligible deterministic contribution of prime powers can be absorbed by using $\rho/2$. Let $a \downarrow 0$ and $b \uparrow 1$ to obtain (10.2).

If $h < \log 2$, this bound is summable on $X = t_0^n$ for every fixed $t_0 > 1$. Intersect over positive rational thresholds to obtain convergence on the dyades for a prescribed family. For one fixed table, use in addition the grids $t_0 = 1 + 1/j$, $j \geq 5$. With $\delta = 1/j$ fixed, changing X by a factor at most $1 + \delta$ changes the positive normalised mass of the cutoffs by $O_{\alpha,\beta}(\delta) + o_\delta(1)$. At a product boundary sum $\Lambda(s)$ over an interval of relative width $O(\delta)$ for each r using the prime number theorem; at a row boundary use partial summation for $\Lambda(r)/r$. Normalisation changes by $O(\delta)$. Send j to infinity on the common event.

At $X = 2^n$ every $F(rs)$ is bounded by a constant times 2^n. Let $J_n = (\log n)/h + O(1)$ be an integer upper bound on the indices of the bits read, chosen within a fixed additive constant of this expression. The event $T_{2^n} > \rho$ is a union of length-J_n cylinders. Each has fair probability

2^{-J_n}, so (10.2) permits at most $n^{o(1)}$ exceptional cylinders, in the sense of $O_\delta(n^\delta)$ for every $\delta > 0$. Their diameters are $O(2^{-J_n})$. For $s > h/\log 2$ the sum over all n of the sth powers of these diameters converges. The limsup cover followed by a countable union over $\rho = 1/j$ gives the upper bound on dimension. The ambient binary space also has dimension one.

For the lower bound put $d_* = \min(1, h/\log 2)$. Free a set of bit positions with exactly $\lfloor d_* j \rfloor$ free positions among the first j, fixing all others to $+1$. The resulting set has dimension d_* by the cylinder count and the uniform measure on its free bits. Each allowed length-j cylinder has mass $2^{-\lfloor d_* j \rfloor}$, giving the matching lower bound by summing the masses of a cover. Across the block $N \le n < 2N$ a common prefix of length $J = \lceil (\log N)/h \rceil + O_h(1)$ suffices. At most $CN^{d_* \log 2/h} \le CN$ admissible prefixes occur. Fixing an additional initial prefix of sufficiently many free bits brings this count down to N on every large block. Cycle through all these initial choices on successive disjoint blocks, repeating each infinitely often. In its block assign every admissible length-J continuation to a separate n. On that dyade prescribe the joint table to be the function $\xi(F(rs))$ computed from the assigned prefix, and fill unused dyades arbitrarily. The construction uses no actual random draw. Every sign sequence in the chosen set agrees with an assigned prefix infinitely often. Its table then squares the sign to one on the entire window. Constant row and column weights give $T_{2^n} \ge \beta - \alpha + o(1)$. When $d_* = 1$ all bit positions are free, so this statement holds for every sign sequence, including at $h = \log 2$. □

C Transferring a twist through the divisor labels

The two cofactors have different sizes. Before removing a twist we need to know how often an exceptional value of a cofactor can occur among the incidences. Throughout this appendix constants may depend on $\boldsymbol{c}$.

Lemma C.1 (The two cofactor projections)**.** *Let $H_m(w)$ and $H_k(w)$ count incidences with the indicated cofactor equal to w. Then*

$$\sum_w H_m(w)^2 \ll A^2 M, \qquad \max_w H_k(w) \ll B.$$

Consequently a set of at most σM possible m's costs $O(X\sqrt{\sigma})$ incidences, and a set of at most σK possible k's costs $O(X\sigma)$.

Proof. For k fixed, $(a,k) = 1$, since a common divisor divides 2 and a is odd. The possible b occupy one residue class modulo a. Summing $O(1 + B/a)$ over $a \asymp A$ gives $O(B)$.

Take $h(w) \ge 0$ supported on the ambient interval $w \asymp M$, and write $e(t) = e^{2\pi i t}$ and $Z(\alpha) = \sum_w h(w) e(w\alpha)$. Dropping the bounds on b and its parity bounds $\sum_w H_m(w)h(w)$ by the sum over $k \asymp K$, $a \asymp A$, $(a,k) = 1$, of the sums of h in the class $-2\bar{a} \bmod k$. Define

$$P_d(r) = \frac{1}{d} \sideset{}{^*}\sum_{j \bmod d} Z(j/d) e(-jr/d).$$

Grouping the additive characters modulo k by their reduced denominator writes the sum in class r exactly as $\sum_{d|k} (d/k) P_d(r)$. For fixed d the coefficients d/k over multiples $k \asymp K$ sum to $O(1)$. Parseval and the additive large sieve [17, Theorem 1] give

$$\sum_{d \asymp D} \sideset{}{^*}\sum_{r \bmod d} |P_d(r)|^2 \ll \frac{M + D^2}{D} \sum_w |h(w)|^2.$$

Indeed distinct reduced fractions here are separated modulo one by $\gg D^{-2}$. For fixed d the class $-2\bar{a} \bmod d$ has $O(1 + A/d)$ preimages among these a, including the possible factor two at even d.

When $D \le A$, Cauchy–Schwarz over $O(D^2)$ pairs (d,r) bounds the block by $O(A(\sqrt{M/D}+\sqrt{D})\|h\|_2)$. When $D > A$, Cauchy–Schwarz over $O(AD)$ pairs (a,d) and bounded multiplicity give $O(\sqrt{A}(\sqrt{M}+D)\|h\|_2)$. The dyadic sums are $O(A\sqrt{M}\|h\|_2)$ because $\sqrt{A}(\sqrt{M}\log X+K) \ll A\sqrt{M}$ for the present exponents. Taking $h = H_m$ proves the first assertion. The size estimates for exceptional sets follow by Cauchy–Schwarz and the maximum bound. □

Proof of Theorem 11.3. For the first site put $\theta_p = \arg u_1(p) \in [-\pi,\pi]$, so $\theta_p^2 \ll |1-u_1(p)|^2$. On $w \le CM$ the contribution of distinct large prime divisors can be replaced by a constant phase using the bound

$$\frac{1}{M}\sum_{w\le CM}\left|\sum_{\substack{p>Y\\ p\mid w}}\theta_p - \sum_{Y<p\le M^{1/4}}\frac{\theta_p}{p}\right|^2 \ll_C E_1(Y)+o(1). \tag{C.1}$$

For the primes up to $M^{1/4}$, count divisibility by p and by pq in $[1,CM]$. Their diagonal variances are bounded by $O_C(\sum\theta_p^2/p)$ and rounding costs $O_C(M^{-1/2})$ in the full second moment around the displayed centre. For larger primes each w has at most five such divisors when X is large. Cauchy–Schwarz and $\lfloor CM/p\rfloor \le CM/p$ bound the remaining mean square by $O_C(E_1(Y))$. The estimate for the other site uses K, E_2.

To use ordinary multiplicativity on $n = am$, exclude the m with some $p^2 \mid m$, $p > Y$. There are $O(M/Y)$ such values. Also exclude $p \mid (a,m)$ for $p > Y$. For $Y < p \le Z = X^{1/40}$ the count of incidences is at most a constant times

$$\sum_{Y<p\le Z}\frac{A}{p}B(1+U/p) \ll X/Y + AB\log\log(3X).$$

For any such a, the cofactor occupies one class modulo b, with $(p,b)=1$, giving the factor $1+U/p$. For $p > Z$ one has $p^2 \mid n$. The number of possible n is $O(X/Z)$ by summing $\lfloor 2X/j^2\rfloor$ over $j > Z$. Each occurs with multiplicity at most $\tau(n)\tau(n+2) = X^{o(1)}$. The same estimates with $b, k, n+2$ apply at the second site. By Lemma C.1, all exclusions together cost $O(X/\sqrt{Y})+o(X)$ incidences at fixed Y.

For the remaining incidences the large-prime part of $u_1(am)/u_1(a)$ equals $\exp(i\sum_{p>Y,\,p\mid m}\theta_p)$, with the cutoff at Y. Equation (C.1), the Lipschitz bound for e^{it}, and Lemma C.1 replace this phase by a constant of modulus one at total cost $O(X\sqrt{E_1(Y)+o(1)})$. The maximum projection gives the analogous replacement at the second site by Cauchy–Schwarz on $k \le CK$. These constants may vary with X.

For $2 < p \le Y$ the remaining factor at the first site is

$$u_1(p^{v_p(a)+v_p(m)})/u_1(p^{v_p(a)}).$$

Cap $v_p(m)$ and $v_p(k)$ at an integer H. The excluded mass is $O(X(\sum_{2<p\le Y}p^{-H})^{1/2})$ for large H, using both projection bounds. A bounded function of a capped valuation is its value at zero plus a sum of successive differences times $\mathbf{1}_{p^j\mid m}$, $1 \le j \le H$. Expand the product over small primes at both sites. The result is a finite linear combination of $\mathbf{1}_{d\mid m}\mathbf{1}_{e\mid k}$ with odd d, e. Its coefficients may depend on a, b and are bounded in terms of Y, H. The prime 2 contributes one by hypothesis. The factors $u_1(a)u_2(b)$ are constant in each cell.

Now $d \mid m, e \mid k$ is equivalent there to $ad \mid n, be \mid n+2$. For fixed d, e the map $(a,b) \mapsto (ad, be)$ is injective. After taking cell moduli the sum is bounded by the original norm on the (d,e)-dilated box times the coefficient bound. For fixed Y, H the assumed cancellation on these boxes makes the finite sum $o(X)$. Take $X \to \infty$, then $H \to \infty$, then $Y \to \infty$. For the reverse implication start with Fu_1, Gu_2 and apply the argument to $\bar{u}_1, \bar{u}_2$. This uses only the modulus bounds on the original F, G. □

If the twists are arbitrary at 2, the same expansion retains filters $2^j \mid m, 2^\ell \mid k$ in both families of norms. Capping and expanding the two valuations at 2 then proves equivalence for that enlarged family. Without these filters the stated odd-dilation argument uses the normalisation in (11.8).

Proof of Proposition 11.4. For either site write $L_0 = M$ or K, $\theta_p = \arg v_{i,X}(p)$. The phase of $v_{i,X}(w)$ is $\sum_{p^j \mid w} \theta_p$ modulo 2π. Centre its part with $p^j \le L_0^{1/4}$ by $\sum_{p^j \le L_0^{1/4}} \theta_p/p^j$. Distinct primes give a total rounding error $O_C(L_0^{-1/2})$ in the mean square on $w \le CL_0$. For powers of the same prime the contribution is bounded using

$$\theta_p^2 \sum_{j,\ell \ge 1} p^{-\max(j,\ell)} \ll \theta_p^2/p.$$

Among the remaining powers with $p > L_0^{1/10}$ an integer w has at most eleven prime factors with multiplicity. Cauchy–Schwarz and the bound $\sum_{j \ge 1} p^{-j} \ll 1/p$ control their mean square by $O_C(\varepsilon_i^2)$. For each $p \le L_0^{1/10}$ the first power above $L_0^{1/4}$ divides w with frequency $O_C(L_0^{-1/4})$. Their union has relative size $O_C(L_0^{-3/20})$, where the error of a unit complex number costs at most two. We obtain a constant phase $\zeta_{i,X}$ with

$$L_0^{-1} \sum_{w \le CL_0} |v_{i,X}(w) - \zeta_{i,X}|^2 \ll_C \varepsilon_i^2 + o(1)$$

uniformly in the prime values. Complete multiplicativity factors out $v_{1,X}(a)v_{2,X}(b)$ exactly. Applying Lemma C.1 to the two errors and using $\big||s| - |t|\big| \le |s - t|$ inside each cell proves the result. □

Proof of Proposition 11.5. In $\lambda(w) = \sum_{r^2 \mid w} \mu(w/r^2)$ exactly one term is nonzero, namely the one with $r = \prod_p p^{\lfloor v_p(w)/2 \rfloor}$. At an odd origin the two square divisors selected this way are odd. Truncating their square roots at R discards $O(M/R)$ possible m and $O(K/R)$ possible k, by summing $\lfloor CM/r^2 \rfloor$ and $\lfloor CK/r^2 \rfloor$ for $r > R$. The incidence cost is $O(X/\sqrt{R})$ by Lemma C.1. For each retained (r, s) replace (a, b) by (ar^2, bs^2). The resulting ordinary Möbius-cofactor cell already vanishes on even origins. Its absolute sums are bounded by the assumed dilated norm. First let X, then R, tend to infinity.

Conversely $\mu(w) = \lambda(w)\mathbf{1}_{w \text{ squarefree}}$. At odd origins approximate this indicator by $\prod_{2<p\le Y}(1 - \mathbf{1}_{p^2 \mid w})$ at incidence cost $O(X/\sqrt{Y})$ for the two sites together. At fixed Y its expansion again dilates the odd labels, with $\lambda(r^2) = 1$, and the assumed odd Liouville norms control each term. Then let $Y \to \infty$. For unrestricted origins, perform both expansions over odd squares only. Their squarefree residue at the odd primes leaves the factor $(-1)^{v_2(w)}$ intact, giving precisely $\mu^{[2]}$. This proves the second equivalence. □

D Conditional means at primes and the fluctuation estimates

All estimates here are uniform in the real array $|v| \le 1$. For $p \in \mathcal{B}$ set $U_p = \{n : p \in T(n)\}$, $D_p = \{n : p \mid P(n)\}$ and $b_p = \mathbb{E}_I[v \mid U_p]$. The band eventually avoids all primes dividing a leading coefficient or a determinant of two forms. Counting their four distinct roots and their lifts gives

$$\mathbb{P}_I(U_p) = \frac{4}{p} + O_P(p^{-2} + X^{-1}), \qquad \mathbb{P}_I(U_p \cap U_s) \ll_P (ps)^{-1} \quad (p \ne s). \tag{D.1}$$

The second inequality uses $Z^2 = X^{o(1)}$. In particular the conditional means are well defined for large X.

A bound for conditional means. The dual form of the Turán–Kubilius inequality compares prime-conditioned means with the initial mean; see Elliott's discussion [7]. In the present short band the CRT gives the constant needed below:

$$4\sum_{p\in\mathcal{B}}\frac{|b_p-M|^2}{p}\le \mathrm{Var}_I(v)+o(1). \tag{D.2}$$

To see this, form the centred vectors $h_p=\sqrt{p/4}\,(\mathbf{1}_{D_p}-\mathbb{P}_I(D_p))$ in $L^2(I)$. Their Gram diagonal is $1-4/p+O_P(p/X)$, and the off-diagonal entries are $O_P(\sqrt{ps}/X)$ by CRT. The absolute error row sums are $O_P(Z^2/X)$, so the squared operator norm of the analysis map $g\mapsto(\langle g,h_p\rangle)_p$ is at most $1+o(1)$. Apply it to $v-M$. Since $\mathbb{P}_I(D_p)=4/p+O_P(1/X)$, this proves (D.2) for conditional means on D_p, up to error $O_P(Z/X)$ in the sum. The removed set $D_p\setminus U_p$ lies in the four classes modulo p^2 above the roots. Replacing each conditional mean changes it by $O_P(1/p+p/X)$, which contributes $o(1)$ to the sum by boundedness. This proves (D.2).

The first Fourier coefficients. Take any subfamily of the band, with independent nonconstant signs of flip probabilities $0<\theta_p\le 1/2$, and take $+1$ elsewhere. Put

$$m_p=1-2\theta_p,\qquad \sigma_p^2=4\theta_p(1-\theta_p),\qquad z_p=(\xi_p-m_p)/\sigma_p,$$
$$J=\sum_p\frac{\sigma_p^2}{p^2},\qquad h=\sum_p\frac{\theta_p}{p},\qquad e_X=R^{-1}+Z/X.$$

All sums in this paragraph run over the active primes. In the orthonormal basis of products of the z_p [19, Chapter 8], the singleton coefficient of $A^\xi(v)$ is

$$c_p=\sigma_p\,\mathbb{E}_I\left[v\,\mathbf{1}_{U_p}\prod_{s\in T(n)\setminus\{p\}}m_s\right].$$

In this product set $m_s=1$ at inactive primes. The cost of dropping the other factors is at most $2\sigma_p\sum_{s\ne p}\theta_s\mathbb{P}_I(U_p\cap U_s)$. By (D.1),

$$\left|c_p-\frac{4\sigma_p}{p}b_p\right|\ll_P(\sigma_p/p)(h+e_X).$$

For the coefficients of degree at least two, take $\mathbb{E}_{\xi_p,\xi_s}[z_pz_sA^\xi(v)]$ at each pair $p<s$, conditioning on the other bits. Its modulus is at most $\sigma_p\sigma_s\mathbb{P}_I(U_p\cap U_s)$. Parseval and summation over pairs count each squared coefficient of degree $j\ge 2$ with multiplicity $\binom{j}{2}$. Consequently their total is $O_P(J^2)$. Together the two estimates give

$$\left\|F^\xi(v)-4\sum_p b_p(\xi_p-m_p)/p\right\|_2\ll_P\sqrt{J}\,(h+e_X+\sqrt{J}). \tag{D.3}$$

Proof of Theorem 12.3. In the first experiment $h\ll 1/\log Z$ by the prime number theorem and $ZJ=H+O(1/\log Z)$. Let $B=\sqrt{Z/H}\sum_p(\xi_p-m_p)/p$, of variance $1+O(1/(H\log Z))$. Since $Z\sigma_p^2/p^2\le 1/p$, (D.2) and orthogonality bound the L^2 distance of $\sqrt{Z/(16H)}F^\xi(v)$ from MB by $O_P(H^{-1/2})$, using (D.3). Taking variances, and taking the covariance with the same quantity for $v=1$, gives the stated estimates. To see $B\to 0$ in probability, split the sum at $\sqrt{Z}$. The chance of any flip below the split is at most $(4Z)^{-1}\sum_{p\le\sqrt{Z}}p\ll 1/\log Z$. When none flips, the corresponding unnormalised centred sum costs $O(1/(\sqrt{Z}\log Z))$. Above the split its variance is $O(Z^{-1}\sum_{\sqrt{Z}<p\le Z}1/p)=O(1/Z)$. This proves (12.4).

For the second experiment write $w = 2^j$. Uniformly for the full intervals, $H_j \asymp 1/\log w$, $K_j \asymp 1/(w \log w)$, and $H' = H + O(1/\log R)$. Here $J = \sigma^2 K_j$, $h = H_j/4$. After division by $4\sigma\sqrt{K_j}$ the error in (D.3) is $O_P(1/\log R) = o((H')^{-1/2})$. The squared distance of its linear term from MG_j, averaged also over j, is

$$\frac{1}{H'} \sum_j \sum_{p \in I_j} \frac{H_j}{p^2 K_j} |b_p - M|^2 \le \frac{2}{H'} \sum_{p \in \mathcal{B}} \frac{|b_p - M|^2}{p}.$$

Indeed $H_j/K_j \le 2w \le 2p$. Taking square roots, adding the linearisation error and using (D.2) proves (12.5). The independent-summand Berry–Esseen inequality [20, Chapter V] bounds the Kolmogorov distance of the conditional law of G_j from a standard Gaussian by

$$C \sum_{p \in I_j} p^{-3}/K_j^{3/2} \ll \sqrt{\log w / w}. \tag{D.4}$$

Each standardised sign has bounded third absolute moment. The bound tends to zero uniformly in j. Equation (12.5) now gives the asserted joint limit. □

Proof of Corollary 12.4. Set $V = G_j(\xi) - G_j(\xi')$ and $U = Y_j(v;\xi) - Y_j(v;\xi')$. Both draws use the same j, so the centres agree and the ratio in (12.6) is U/V when $V \neq 0$. By (12.5), $\mathbb{E}|U - MV|^2 \ll_P 1/H'$. Given j, the summands of $V/\sqrt{2}$ are independent with total variance one and the same third-moment bound as in (D.4), up to an absolute factor. Thus $\mathbb{P}(|V| \le s) \ll s + \sqrt{\log(2R)/R}$, including the possibility of a zero denominator. On $|V| > s$, Chebyshev bounds the probability of a ratio error exceeding ε by $O_P((H'\varepsilon^2 s^2)^{-1})$. Take $s = (H'\varepsilon^2)^{-1/3}$. This argument needs no convergence hypothesis on M. □

Assistance and supplementary material. GPT-6 Astra assisted with the research and writing on prime clocks and polynomial traces. OpenAI Codex assisted with the prime-pair and cofactor arguments, the periodic and fluctuation estimates, bibliography, exposition and finite checks. The Zenodo source package includes executable diagnostics for the finite identities. The author retains responsibility for the manuscript.

References

[1] S. Arango-Piñeros, D. Keliher and C. Keyes, *Mertens' theorem for Chebotarev sets*, International Journal of Number Theory **18** (2022), no. 8, 1823–1842. doi: 10.1142/S1793042122500932. arXiv:2103.14747.

[2] M. Aymone and V. Sidoravicius, *Partial sums of biased random multiplicative functions*, Journal of Number Theory **172** (2017), 343–382. doi: 10.1016/j.jnt.2016.08.020. arXiv:1412.1157.

[3] D. Berend and Y. Bilu, *Polynomials with roots modulo every integer*, Proceedings of the American Mathematical Society **124** (1996), no. 6, 1663–1671. doi: 10.1090/S0002-9939-96-03210-8.

[4] E. Bombieri, *On the large sieve*, Mathematika **12** (1965), 201–225. doi: 10.1112/S0025579300005313.

[5] R. de la Bretèche, V. Y. Wang and M. W. Xu, *Random Multiplicative Functions and Making Squares from Polynomial Values*, preprint, 2026. arXiv:2607.06398v1. doi: 10.48550/arXiv.2607.06398. Theorem numbering refers to this version.

[6] B. Efron and C. Stein, *The jackknife estimate of variance*, The Annals of Statistics **9** (1981), no. 3, 586–596. doi: 10.1214/aos/1176345462.

[7] P. D. T. A. Elliott, *High-power analogues of the Turán–Kubilius inequality, and an application to number theory*, Canadian Journal of Mathematics **32** (1980), no. 4, 893–907. doi: 10.4153/CJM-1980-068-0.

[8] C. Elsholtz, B. Klahn and M. Technau, *On polynomials with roots modulo almost all primes*, Acta Arithmetica **205** (2022), no. 3, 251–263. doi: 10.4064/aa220407-9-7. arXiv:2206.13466v2.

[9] A. Fish, *Random Liouville functions and normal sets*, Acta Arithmetica **120** (2005), no. 2, 191–196. doi: 10.4064/aa120-2-6. arXiv:math/0509315.

[10] K. Ford, *Poisson approximation of prime divisors of shifted primes*, International Mathematics Research Notices **2025** (2025), no. 7, rnaf079. doi: 10.1093/imrn/rnaf079. arXiv:2408.03803.

[11] S. R. Garcia, E. S. Lee, J. Suh and J. Yu, *An effective analytic formula for the number of distinct irreducible factors of a polynomial*, Journal of the Australian Mathematical Society **113** (2022), no. 3, 339–356. doi: 10.1017/S1446788721000227. arXiv:2012.03102.

[12] A. J. Harper, *Moments of random multiplicative functions, I: Low moments, better than squareroot cancellation, and critical multiplicative chaos*, Forum of Mathematics, Pi **8** (2020), article e1. doi: 10.1017/fmp.2019.7.

[13] H. Iwaniec and E. Kowalski, *Analytic Number Theory*, American Mathematical Society Colloquium Publications **53**, Providence, RI, 2004. doi: 10.1090/coll/053.

[14] O. Klurman, A. P. Mangerel and J. Teräväinen, *On Elliott's conjecture and applications*, Proceedings of the London Mathematical Society **133** (2026), no. 3, article e70205. doi: 10.1112/plms.70205. Theorem and corollary numbers cited here refer to arXiv:2304.05344v2; editorial numbering is not used.

[15] O. Klurman, A. P. Mangerel and J. Teräväinen, *Multiplicative functions in short arithmetic progressions*, Proceedings of the London Mathematical Society **127** (2023), no. 2, 366–446. doi: 10.1112/plms.12546.

[16] C. McDiarmid, *On the method of bounded differences*, in J. Siemons (ed.), *Surveys in Combinatorics, 1989*, London Mathematical Society Lecture Note Series **141**, Cambridge University Press, 1989, 148–188. doi: 10.1017/CBO9781107359949.008.

[17] H. L. Montgomery and R. C. Vaughan, *The large sieve*, Mathematika **20** (1973), 119–134. doi: 10.1112/S0025579300004708.

[18] J. Najnudel, *On consecutive values of random completely multiplicative functions*, Electronic Journal of Probability **25** (2020), article 59, 28 pp. doi: 10.1214/20-EJP456. arXiv:1702.01470.

[19] R. O'Donnell, *Analysis of Boolean Functions*, Cambridge University Press, 2014. Corrected text (2021), arXiv:2105.10386. doi: 10.1017/CBO9781139814782.

[20] V. V. Petrov, *Sums of Independent Random Variables*, Ergebnisse der Mathematik und ihrer Grenzgebiete, 2. Folge, **82**, Springer, Berlin, 1975, Chapter V, pp. 104–133. doi: 10.1007/978-3-642-65809-9.

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

[22] J.-P. Serre, *On a theorem of Jordan*, Bulletin of the American Mathematical Society **40** (2003), no. 4, 429–440. doi: 10.1090/S0273-0979-03-00992-3.

[23] K. Soundararajan and M. W. Xu, *Central limit theorems for random multiplicative functions*, Journal d'Analyse Mathématique **151** (2023), 343–374. doi: 10.1007/s11854-023-0331-y.

[24] J. M. Steele, *An Efron–Stein inequality for nonsymmetric statistics*, The Annals of Statistics **14** (1986), no. 2, 753–758. doi: 10.1214/aos/1176349952.

[25] L. Warnke, *On the method of typical bounded differences*, Combinatorics, Probability and Computing **25** (2016), no. 2, 269–299. doi: 10.1017/S0963548315000103. arXiv:1212.5796. Only the classical bounded-differences inequality recalled in Theorem 1 is used.