hexagon:2611.00798v1

HOOLEY'S TWISTED CANCELLATION BELOW LEVEL 4/15 DOES NOT FORCE PRIMES

WREN AQUILA

Abstract. We work in Ford and Maynard's framework for sieve problems with Type I and Type II information. To Type I information at level 1/2 and Type II information for factors up to x^ν, we add Hooley's condition: the parity bias of the sequence in arithmetic progressions is small on average over moduli up to x^α. We construct prime-free sequences that satisfy all three for every Type II range $\nu < 1/30$ and every twisted level $\alpha < 4/15$, and strictly below each of four further corners (η, α') of a staircase in the (ν, α)-plane, from $(1/11, 1/10)$ to $(1/20, 1/5)$, that is, for every $\nu < \eta$ and every $\alpha < \alpha'$. The constructions add to Ford and Maynard's continuum model one function of the total size of the small prime factors for each number of them; these functions solve the twisted conditions exactly. Each case rests on a computer-assisted certificate in exact or directed-rounding integer arithmetic. Independent audits read the certificates' code for every case except $(1/11, 1/10)$ and $(1/14, 1/8)$, whose certificates had an independent review of their formulas and inputs but not of their code. The cases $(1/11, 1/10)$, $(1/14, 1/8)$ and $(1/30, 4/15)$ each rest on a single implementation. Conversely, the term $d = 1$ of Hooley's condition, global parity balance, already forces a positive proportion of primes once $\nu \geq 0.116$. For $\alpha > 1/3$ we exhibit twisted rows that the added functions cannot enter, and we conjecture, on floating-point evidence, that constructions of exactly this form cannot pass $\alpha = 1/3$ (Conjecture 8.2). We also record what the constructions imply for sieve approaches to Goldbach's problem.

1. Introduction

1.1. **Type I, Type II and parity.** Let (a_n) be a nonnegative sequence supported on $(x/2, x]$, compared with a sequence (b_n) whose prime sum is understood. A sieve method bounds $\sum_p a_p$ in terms of $\sum_p b_p$, using only how well $w = a - b$ is distributed: in arithmetic progressions (Type I information) and against bilinear forms (Type II information). Ford and Maynard [5] recently set this up as a general framework. Type I information (I) holds up to level x^γ, and Type II information (II) holds for one factor in $((x/2)^\theta, x^{\theta+\nu}]$. They define the best constants $C^-(\gamma, \theta, \nu) \leq C^+(\gamma, \theta, \nu)$ with $(C^- + o(1)) \sum_p b_p \leq \sum_p a_p \leq (C^+ + o(1)) \sum_p b_p$ [5, Definition 4.8]; §2 recalls the definitions.

Type II information is indispensable. With $\nu = 0$, Selberg's sequence $a_n = 1 + \lambda(n)$, with $b_n = 1$ and λ the Liouville function, satisfies (I) for every $\gamma < 1$ and has no primes [17], [5, §1]. This is the parity phenomenon. Ford and Maynard show that a small Type II range does not suffice either: for every $\gamma < 1$ there is $\nu_0(\gamma) > 0$ such that $C^-(\gamma, \theta, \nu) = 0$ for $\nu \leq \nu_0(\gamma)$, which they describe as "the first examples (at least with $\gamma \geqslant 1/2$ and $\nu > 0$)" of prime-free sequences with nontrivial Type I and Type II information [5, Theorem 2.1 and §2.1]. In the classical case $\gamma = 1/2$, $\theta = 0$, Ford and Maynard state that primes are forced at $\nu = 0.16623$ and for every $\nu \geq 1/5$, and show $C^-(1/2, 0, 0.1616) = 0$ [5, Theorems 2.2 and 2.7, the proofs of 2.7(b), (c), and the case $\nu \geq 1 - \gamma$ of §10.1] (for $\nu \geq 1/3$ see also §7.5). Ford has since reported a prime-free example at $\nu = 0.163$ [4, p. 2728].

1.2. **Hooley's condition.** Selberg's example is extremely biased in its parity: $\sum_n \mu(n) a_n$ is of the order of x. Under the density, growth and squarefree-support hypotheses of their Theorem 1, Friedlander and Iwaniec combine Type I information at a level $x^{2/3} < D < x$ with their Möbius bilinear condition to obtain the prime-sum formula with relative error $O(\log \delta / \log \Delta)$. It is an

Date: 21 November 2026.

asymptotic when their bilinear-range parameters satisfy $\log\delta = o(\log\Delta)$ [6, Theorem 1]. Ford asked whether, for each Type I level $\vartheta < 1$, there are sequences with Type I information at level ϑ and no global parity bias, $\sum_{n\le x}\mu(n)a_n \ll_B x(\log x)^{-B}$, for which Bombieri's asymptotic formulas $\sum_{n\le x} a_n\Lambda_k(n) \sim kx(\log x)^{k-1}$, with $\Lambda_k = \mu * \log^k$, fail. He constructed such sequences for which they fail for every $k \ge 2$ and hold for $k = 1$, the case of primes [3, p. 1665 and Theorem 2]. He called it an interesting problem to replace global parity balance by a stronger condition, but one weaker than Friedlander and Iwaniec's, and as one possibility recorded Hooley's suggestion "that the parity bias in arithmetic progressions is small on average", in the form

$$\sum_{d\le x^{\alpha}}\Big|\sum_{\substack{n\le x\\ d\mid n}}\mu(n)a_n\Big| \ll_B \frac{x}{(\log x)^B}$$

[3, p. 1665]. We write (H_α) for this condition, imposed on $w = a - b$ over $(x/2, x]$, and $C^-(\gamma,\theta,\nu \mid \mathrm{H}_\alpha)$ for Ford and Maynard's constant over the sequences that also satisfy it (§2). The term $d = 1$ alone is the global parity condition (Par), a dyadic form of Ford's condition. Ford's question involves no Type II information; this paper studies what (Par) and (H_α) do when Type II information is added. Ford and Maynard also ask for "a theory which can incorporate such additional arithmetic information" (their examples are trilinear estimates, density assumptions and estimates for special coefficients), and add that they "have not attempted such a generalization" [5, §4.6].

Twisted information of related kinds is known to break the parity barrier. For Goldbach's problem, Huang and Li fix $A > 0$ and $0 < \theta < 1$, and assume the Elliott–Halberstam conjecture for the primes at level $N^\theta(\log N)^{2A+8}$ together with its variant twisted by the Möbius function at level $N^{1-\theta}$: for every $B > 0$,

$$\sum_{q\le N^{1-\theta}}\max_{y<N}\max_{(a,q)=1}\Big|\sum_{\substack{n\le y\\ n\equiv a \bmod q}}\Lambda(n)\mu(N-n) - \frac{1}{\varphi(q)}\sum_{n\le y}\Lambda(n)\mu(N-n)\Big| \ll_B \frac{N}{(\log N)^B}.$$

Under these hypotheses $\sum_{n<N}\Lambda(n)\Lambda(N-n) \ge \mathfrak{S}(N)(1-\mathcal{A}(N))N + O(N(\log N)^{-A})$ for every even N, where $\mathfrak{S}(N)$ is the singular series and $\mathcal{A}(N) = \prod_{p\nmid N,\, p>2}\big(1-\frac{1}{p(p-1)}\big)$ [11, Theorem 1]. They deduce Goldbach's conjecture for all large even N from the same hypotheses with the first at level $N^\theta(\log N)^{12+\epsilon}$, and, since the Bombieri–Vinogradov theorem supplies that hypothesis for every $\theta < 1/2$, from the twisted hypothesis at some level above $1/2$ alone [11, Corollary 1]. Moya [15] reports that the divisor switch in [11, (18)] drops a moving endpoint, $n < N - \alpha k$ with α their divisor cut-off, and states that the conclusion holds under the original hypotheses with the endpoint restored; this is one reason the interval maximum matters. For a fixed nonzero even integer h, Murty and Vatwani assume $EH_\Lambda(x^\theta(\log x)^C)$ and $EH_{\mu_h}(x^{1-\theta})$ for a fixed $\theta < 1$ and a sufficiently large fixed C, where the second hypothesis is their centred progression estimate for $\Lambda(n)\mu(n+h)$. They prove $\sum_{n\le x}\Lambda(n)\Lambda(n+h) \ge (1-o(1))\mathfrak{S}(h)(1-\mathcal{A}(h))x$, with $\mathcal{A}(h) = \prod_{p\nmid h,\ p>2}(1-1/(p(p-1)))$ [16, Theorem 1.1]; only positive arguments of the arithmetic functions are included. These hypotheses concern particular sequences. They subtract the full initial-interval sum divided by $\varphi(q)$ and take maxima over reduced residue classes and initial intervals. They are not the condition (H_α) for a general sequence. For general sequences, Iwaniec's sieve for primes and Friedlander and Iwaniec's inequality combine Type I information at level $x^{1-\delta}$ with a bilinear form in which the Möbius function weights a factor in $(x^\delta, x^{2\delta}]$ [12, §6, Proposition, p. 296], [8, Proposition 6.1] (a variant of [7, Theorem 18.5]); Vatwani treats twin primes under an Elliott–Halberstam-type conjecture for the primes p with $p+2$ squarefree, at every level below 1, together with a Chowla–Elliott-type bound $\sum_{d<x^\theta}|\sum_{p\le x,\ d\mid p+2}\mu(p+2)| \ll x(\log x)^{-2}$ for every $\theta < 1/2$, which has the shape of (H_θ) for that particular sequence [18, Chapter 7]. Whether (H_α) at a level $\alpha > 1/2$ forces primes for general sequences in Ford and Maynard's framework is not addressed in this paper beyond Theorem B, by which, with $\gamma = 1/2$ and $\theta = 0$, the term $d = 1$ of (H_α) alone already forces primes once $\nu \ge 0.116$.

In the plane of the Type II range ν and the twisted level α, with Type I information at level $1/2$, Ford and Maynard state results that force primes at every α for $\nu = 0.16623$ and for every $\nu \geq 1/5$ [5, Theorems 2.2 and 2.7(a), the proof of 2.7(b), and the case $\nu \geq 1-\gamma$ of §10.1] (for $\nu \geq 1/3$ see also §7.5). This paper asks how far into the rest of this plane prime-free sequences reach.

1.3. **Results.** Our first theorem gives prime-free sequences with Type I information at level $1/2$, a positive Type II range, and Hooley's condition at every twisted level below $4/15$.

Theorem A (The ladder)**.** *Let (η, α') be one of*

$$(1/11,\ 1/10), \quad (1/14,\ 1/8), \quad (1/18,\ 1/6), \quad (1/20,\ 1/5), \quad (1/30,\ 4/15).$$

Fix $0 < \nu < \eta$ and $0 \leq \alpha < \alpha'$. There is a uniformly bounded family of sequences $a_n = a_n^{(x)} \geq 0$ supported on $(x/2, x]$ with the following property. For every $B > 0$ and every sufficiently large x, the sequence $w_n = a_n - 1$ satisfies (I) with $\gamma = 1/2$, (II) with $\theta = 0$ on $1 < \ell \leq x^\nu$, and (H_α). Moreover, $a_p = 0$ for every prime $p \in (x/2, x]$. In particular, $C^-(1/2, 0, \nu \mid \mathrm{H}_\alpha) = 0$.

The five cases are the corners of a staircase in the (ν, α)-plane. In our constructions a larger twisted level came at the cost of a shorter Type II range; we do not know whether this trade-off is forced. The case $(1/30, 1/4)$, contained in $(1/30, 4/15)$, is proved on the way; it is the input of one of the Goldbach consequences in §9. The simplest construction, $(\eta, \alpha') = (1/11, 1/11)$, is implied by the case $(1/11, 1/10)$ and is described in Remark 6.1 for illustration.

Our second theorem says that the term $d = 1$ alone, global parity balance, already forces primes once $\nu \geq 0.116$. Without it, Ford and Maynard's examples give $C^-(1/2, 0, \nu) = 0$ for every $\nu \leq 0.1616$ [5, Theorem 2.7(c)]: their construction at $\nu = 0.16169$ satisfies (I) and (II) there, hence at every smaller ν, where the Type II range only shrinks, and its $b_n = 1$ and bounded a_n give the remaining conditions [5, Theorem 6.3(b) and Lemma 4.6]. So on $[0.116, 0.1616]$ the parity condition changes the answer.

Theorem B (Parity forces primes from $\nu = 0.116$)**.** (a) *For $0.116 \leq \nu < 1/3$,*

$$C^-(1/2, 0, \nu \mid \mathrm{Par}) \ \geq\ 1 - 2I_4(\nu) - 4I_5(\nu) - \tfrac{20}{3} I_6(\nu) - \tfrac{25}{3} I_7(\nu) - \tfrac{35}{3} I_8(\nu) \ >\ 0,$$

where $I_k(\nu) = \frac{1}{k!} \int \frac{dx_1 \cdots dx_{k-1}}{x_1 \cdots x_k}$ over $x_i \in (\nu, 1-\nu)$ with $x_1 + \cdots + x_k = 1$, and $I_k = 0$ if $k\nu \geq 1$; in particular $I_8 = 0$ for $\nu \geq 1/8$. The bound is at least 0.0350 throughout the range, at least 0.2892 for $\nu \geq 1/8$, at least 0.6198 for $\nu \geq 1/7$, and at least 0.7072 for $\nu \geq 3/20$.

(b) *For $1/4 < \nu < 1/3$, (w), (I), (II) and (Par) give $\sum_{x/2 < p \leq x} (a_p - b_p) \ll_A x(\log x)^{-A}$ once B is large enough in terms of A, ϖ and ν; in particular $C^-(1/2, 0, \nu \mid \mathrm{Par}) = 1$.*

Moreover, $C^-(1/2, 0, \nu \mid \mathrm{H}_\alpha) > 0$ for every $0 \leq \alpha < 1$ and every $0.116 \leq \nu \leq 1$: for $\nu < 1/3$ this follows from (a), since the term $d = 1$ of (H_α) is (Par); for $1/3 \leq \nu \leq 1$, (w), (I) and (II) with $\gamma = 1/2$ and $\theta = 0$ already give $\sum_{x/2 < p \leq x} (a_p - b_p) \ll_A x(\log x)^{-A}$ once B is large enough (§7.5; for $1/3 \leq \nu \leq 1/2$ Ford and Maynard state this asymptotic in [5, Theorem 2.2(a)]*), so that $C^-(1/2, 0, \nu) \geq 1$ by (b.1), and hence $C^-(1/2, 0, \nu \mid \mathrm{H}_\alpha) \geq 1$ (§2).*

So the region where twisted information at Type I level $1/2$ might force primes lies, as far as we know, between the staircase of Theorem A and the line $\nu = 0.116$. Neither theorem says anything about it. The line $\nu = 0.116$ is where the computation reported here stops; at $\nu = 0.116$ the right-hand side of Theorem B(a) is still at least 0.0350. In floating point, that right-hand side changes sign between $\nu = 0.114$ and 0.115.

The proofs of Theorem B(a) and (b) use Ford and Maynard's arithmetic estimate [5, (7.15)], which we prove in full for $\gamma \geq 1/2$ (Lemma 7.1), with corrections to its printed proof, one of them at repeated prime factors: their identification of the ordinary-divisor weight with a vector convolution [5, Lemma 7.18(a)] can fail at non-squarefree n, and we keep the two weights apart (§7). The final assertion of Theorem B for $\nu \geq 1/3$ does not use [5, (7.15)]; §7.5 proves it directly.

1.4. **How the results are established.** Every case of Theorem A rests on a computer-assisted certificate: a rigorous enclosure, in exact integer or directed-rounding arithmetic, showing that an explicit finite-dimensional function satisfies strict inequalities with a positive margin (§5). For $(1/18, 1/6)$ and $(1/20, 1/5)$, independent audits read the code bodies and reran the computations in full at reduced resolution; the second-order runs of these two cases were also rerun at their own resolution, where they pass, and the separate $(1/20, 1/5)$ implementation was replayed by an audit only at $M = 36{,}000$, where it does not pass. For $(1/30, 1/4)$ the audit read the code bodies and reran the computation in full at $M = 1920$, where the margin is negative: the rerun checks the code path and that its enclosures contain those of the stored $M = 7680$ output, not the certified margin; the run of record was not rerun (§5.5). For $(1/30, 4/15)$ the audits read the code bodies and reran parts of the computation: the forcing chain, the parity solve, $f(1)$, the free bound and the adaptively refined extended class at $M = 960$, and the diagonals and the parity data at the full resolution $M = 23{,}040$. No audit reran the full computation at any resolution, and none recomputed the extended class that binds its margin (§5). For $(1/11, 1/10)$ and $(1/14, 1/8)$ the certificates' code was not audited: an independent review checked their formulas, inputs, signs, normalisations and selected parts of the exact arithmetic against the text, and recomputed in exact arithmetic the least-margin cell of $(1/11, 1/10)$, using that certificate's enclosures of h and $f(1)$. Neither certificate was rerun. The cases differ in how many independent certificates exist.

- $(1/18, 1/6)$ and $(1/20, 1/5)$: certificates from two separate implementations each, in the sense of §5.6. For $(1/18, 1/6)$ these are its first-order exact-integer certificate and the separate implementation's; its third certificate is a run of the second-order engine, which was developed from the first-order code and is not counted as separate.
- $(1/14, 1/8)$ and $(1/11, 1/10)$: one exact-integer certificate each. A second computation, in floating point with outward rounding, agrees, but its error control rests on an a priori round-off bound for library FFT convolutions, which we do not count as rigorous.
- $(1/30, 4/15)$, and the contained case $(1/30, 1/4)$: one implementation each, with independent audits of its code (§5.5 states what they checked). No second certificate exists. For $(1/30, 1/4)$ a second, independent first-order certificate was implemented; measured at small resolution on its own code, it was projected to need M of 4 to 13 million for a positive margin, and the full run was not attempted. On that projection a second first-order certificate is out of reach at affordable cost. A separately written higher-order implementation is the natural route to a second certificate; none exists. For $(1/30, 4/15)$ no second certificate was attempted.

The numerical bounds of Theorem B rest on exact-integer certificates computed by a separately written program during the audit of that theorem and, below $\nu = 1/8$, by a copy of that program that adds the term in I_8 to the bound and leaves its enclosure arithmetic unchanged; the copy was run after the audit, and with that term set to zero it reproduces the audit's stored exact outputs at $\nu = 3/20$, $1/7$ and $1/8$ byte for byte. At $\nu = 0.116$ a second, separately written program, which applies the same cell-convolution scheme in exact integer arithmetic on a mesh of step $1/1{,}024{,}000$, gives a lower bound of at least 0.0354; it was also written and run after the original audit, which covered neither new program. No exact computation of the bounds at $\nu = 3/20$, $1/7$ and $1/8$ by a second, separately written program is reported here. These three programs and their outputs are in the supplement's `thmB/` directory.

1.5. **Method.** The sequences are Ford and Maynard's: $a_n = 1 + F(v(n))$ on squarefree η-rough n, where $v(n)$ is the vector of $\log p / \log n$ over the prime factors p of n, and F is a function on such vectors that satisfies their fragmentation relation (6.3), so that (I) at level $1/2$ holds automatically [5, §6]. Type II information on $1 < \ell \le x^\nu$ is automatic for $\nu < \eta$: such an ℓ has a prime factor at most x^ν, which is less than $(x/2)^\eta$ for large x, so ℓ divides no element of the support of the perturbation $a_n - 1$, whose prime factors are all at least $n^\eta > (x/2)^\eta$. Ford and Maynard's

examples for their Theorem 2.1 are built from their modified Liouville functions, products over the components of a one-variable function [5, §§9.1–9.2]. The main term is $(-1)^k$ on vectors whose k components all lie below $1-\gamma$; a correction term also depends on the sizes of the components. These examples are built for Type I and Type II information and have no reason to satisfy twisted conditions. We add, for each number j of "small" components (those in $[\eta, \alpha']$), a function h_j of their total size. The twisted rows at moduli $p_1 \cdots p_t$ are then linear in h_t with a nonvanishing diagonal coefficient, so they can be solved exactly, top-down in t (§3). A linear program chooses the remaining constants, and a rigorous enclosure certifies, with a margin, the inequalities that make $a_n \geq 0$ once f is normalised so that $a_p = 0$. Ford and Maynard's construction procedure (the proof of [5, Theorem 6.3]), with their mask, which makes the function vanish on every vector with a subsum near $1/2$ or just above ν (the vector $E = (1/2, 1/2)$, where the certificate asserts no bound, among them), followed by a re-solve of the forced functions and a normalisation, then turns the continuum function into sequences with $a_p = 0$ exactly (§4).

1.6. **Limits.** The construction meets an obstruction past $1/3$. For $\alpha > 1/3$, some twisted rows involve moduli that no vector of the support can complete with rough components alone, and we prove that the forced functions h_j do not enter these rows (Lemma 8.1). Floating-point experiments (for the profile $w_k = 1$ at $\eta = 1/36$, on a grid of step $1/72$) suggest that the method cannot pass $1/3$, but that part is a conjecture about this family of constructions (Conjecture 8.2), not a statement about Hooley's condition: on three of the four grids we could compute, a discretised version of the problem with fully free data, capped at 4 times the baseline and with Type I information that includes moduli of size exactly $x^{1/2}$, does get past $1/3$ (with exact certificates on two of them). How close to $1/3$ the family can get is not known: the certified cases reach $4/15$, and in floating point, on one grid, the family's values just past $1/3$ stay far above -1 for every box and bin choice we tried, while for $\eta \leq 1/36$ richer data that also see the largest small component come within 0.05 of -1 at the least level computed at or past $1/3$ on each grid (below it on two of the four grids under both boundary conventions tested, with the other two unresolved under the tested conventions and boxes), from which we draw no conclusion.

1.7. **Consequences for Goldbach.** The constructions reflect to Goldbach's problem (§9). For each case of Table 4 and every large even N there is a nonnegative sequence on $[1, N)$ that vanishes on every prime $m \geq N/(\log N)^C$, and that agrees with a presieved local model of $\Lambda(N-m)$ in five respects: Type I sums at the Bombieri–Vinogradov level, small-factor Type II sums (exactly), Hooley's twisted condition at level α, parity sums over every interval, and roughness comparisons up to $(N/(2(\log N)^C))^\eta$. So these inputs do not by themselves force a Goldbach-scale count, even with Hooley's condition granted up to any fixed level $\alpha < 4/15$ (with Type II below $N^{1/30}$). The Type II identity, the twisted condition and the parity sums at the precision $N(\log N)^{-A}$ are properties of the reflected sequence. In the arbitrary-coefficient form the reflection uses, the true sequence fails the first, as it fails the exact roughness identity, and it is not known to have the other two. For it, even $\sum_{p<N} \mu(N-p) = o(\pi(N))$ has been obtained only recently and conditionally [9], building on Lü, Wang and Wu's work on $\sum_{p \leq x} \mu(p+2)$ [14]. §9 says what the audits of the reflection checked, and each case carries the certification status of its construction (§5.6).

1.8. **Prior work and credit.** The continuum model, of vectors of prime factors and Type I-compatible functions [5, Definitions 5.1 and 6.2], the fragmentation relation [5, (6.3), Theorem 6.4], the mask and the construction theorem [5, Theorem 6.3 and its proof in §6.2] are Ford and Maynard's. Ford's constructions for Bombieri's asymptotic sieve [2] rest on a related observation: his perturbations $a_n - 1$ are supported on integers whose prime factors lie in windows near $x^{i/M}$, and moduli outside these windows see only the baseline [3, (3.2) and p. 1669; p. 1673]. Ford remarks that the sequences constructed for his Theorem 1 satisfy Hooley's condition for every fixed $0 \leq \alpha < 1-\vartheta$, where ϑ (his ν) is the Type I level [3, pp. 1663, 1665–1666]. This does not hold as printed, and we

do not use it: for those sequences, $\left|\sum_{n\le x}\mu(n)a_n\right| \gg x/\log x$ for arbitrarily large x, so the $d=1$ term already violates the condition [3, p. 1667; (3.9)–(3.10), pp. 1671–1672; p. 1673]; Ford notes at his (1.5) that these sequences "all exhibit a 'global parity bias'" [3, p. 1665]. The lower bound in Theorem B(a) follows the proof of Ford and Maynard's sieve bound [5, Theorem 7.3], with one additional parity term and a separate sign condition at non-squarefree integers. We are not aware of an earlier prime-free construction satisfying the twisted condition, including its $d=1$ term.

1.9. **Organisation.** §2 sets up the framework and the continuum model. §3 gives the construction and §4 the passage to sequences. §5 describes the certificates and their audits, and §6 the individual cases. §7 proves Theorem B. §8 discusses the obstruction past 1/3 and Conjecture 8.2, and §9 the consequences for Goldbach. §10 and the supplement describe how to reproduce the computations.

2. Setting

2.1. **Type I, Type II and Hooley's condition.** We follow Ford and Maynard [5]. Let (a_n) and (b_n) be real sequences supported on $(x/2,x]$, with $a_n\ge 0$, and put $w_n=a_n-b_n$. For parameters $0<\gamma<1$, $0\le\theta<1/2$ and $0<\nu\le 1-\theta$, a constant $\varpi>1$ and a large constant B, the hypotheses are:

(I) (Type I range $[0,1-\gamma]$) $\displaystyle\sum_{m\le x^\gamma}\tau^B(m)\max_I\Big|\sum_{\substack{x/2<mn\le x\\ n\in I}} w_{mn}\Big| \le \frac{x}{\log^B x}$, the maximum over intervals I;

(II) (Type II range $[\theta,\theta+\nu]$) $\displaystyle\Big|\sum_{\substack{(x/2)^\theta<\ell\le x^{\theta+\nu}\\ x/2<\ell n\le x}}\xi_\ell\kappa_n w_{\ell n}\Big| \le \frac{x}{\log^B x}$ for all $|\xi_\ell|\le\tau^B(\ell)$, $|\kappa_n|\le\tau^B(n)$;

(w) (Ford and Maynard's mild growth condition [5, §2]) $\displaystyle\sum_n |w_n|\tau(n)\le x(\log x)^{\varpi}$ and $w_n\ge -x^{\nu/10}$ for $x/2<n\le x$.

FM also require $b_n\ge 0$ and two conditions on the comparison sequence: a lower bound $\sum_p b_p\ge x/(\log x)^{\varpi}$, their (b.1), and a prime number theorem for products of primes weighted by b_n, their (b.2) [5, §4.2]. For given $\varpi>1$ and B, the pairs $((a_n),(b_n))$ with all these properties form their class Ψ [5, Definition 4.7]. $C^-(\gamma,\theta,\nu)$ is the supremum of the constants C such that for every $\varpi>1$ there is B with $\sum_p a_p\ge C\sum_p b_p$ for every pair in Ψ and every sufficiently large x, and $C^+(\gamma,\theta,\nu)$ is defined symmetrically [5, Definition 4.8].

For $0\le\alpha<1$ we add *Hooley's condition at level* α. Hooley suggested postulating "that the parity bias in arithmetic progressions is small on average", and Ford records it in the form $\sum_{d\le x^\alpha}\left|\sum_{n\le x,\,d\mid n}\mu(n)a_n\right| \ll_B x(\log x)^{-B}$ [3, p. 1665]. Since FM's comparison sequence b_n is general, we impose the condition on $w=a-b$:

$$(\mathrm{H}_\alpha)\qquad \sum_{d\le x^\alpha}\Big|\sum_{\substack{x/2<n\le x\\ d\mid n}}\mu(n)\,w_n\Big| \le \frac{x}{\log^B x}.$$

The term $d=1$ is the global parity condition

$$(\mathrm{Par})\qquad \Big|\sum_{x/2<n\le x}\mu(n)\,w_n\Big| \le \frac{x}{\log^B x},$$

so (H_0) is (Par). We write $C^{\pm}(\gamma,\theta,\nu\mid \mathrm{H}_\alpha)$ for the constants of [5, Definitions 4.7 and 4.8] with (H_α) adjoined to (I) and (II), with FM's order of quantifiers (B large, then x large). Adjoining a condition can only shrink the class of sequences, so $C^-(\gamma,\theta,\nu\mid \mathrm{H}_\alpha)\ge C^-(\gamma,\theta,\nu\mid \mathrm{Par})\ge C^-(\gamma,\theta,\nu)$ for every $0\le\alpha<1$. All our results have $\gamma=1/2$ and $\theta=0$. Our prime-free constructions use $b_n=1$ on $(x/2,x]$: we give a family $a_n=a_n^{(x)}\ge 0$, bounded uniformly in x and independent of B, that for

every $B > 0$ and every sufficiently large x satisfies (I), (II) on $1 < \ell \le x^\nu$, and (H_α), with $a_p = 0$ at every prime in $(x/2, x]$. Condition (w) holds automatically for these bounded families, and (b.1) and (b.2) hold for $b_n = 1$ by [5, Lemma 4.6], with $q = 1$, $y = x/2$. Thus $C^-(1/2, 0, \nu \mid \mathrm{H}_\alpha) = 0$ for the parameters of these constructions. Theorem B instead applies to general comparison sequences satisfying (b.1) and (b.2), and for $\nu < 1/3$ uses a sieve argument with an additional parity term.

With $b_n = 1$, Hooley's form for a_n and (H_α) for w_n are equivalent for every fixed $\alpha < 1$, because the baseline sums are small. Write $n = dm$. The inner sum for b is $\mu(d)$ times a difference of two sums $\sum_{m \le y,\,(m,d)=1} \mu(m) = \sum_{k \le y,\ k \mid d^\infty} M(y/k)$, with M Mertens' function and $y = x/d$, $x/(2d)$. Split at $k = \sqrt{y}$; use $|M(y)| \ll y\, e^{-c\sqrt{\log y}}$ below the split and Rankin's trick above it. The total over $d \le x^\alpha$ is $\ll x \log x \log\log x\; e^{-c\sqrt{\log x}} + x^{3/4+\alpha/4+o(1)}$.

2.2. **The continuum model.** Let $0 < \eta < 1/2$. A squarefree integer n is *η-rough* if every prime factor of n is at least n^η. For a modulus $d \le x^\alpha$ we use the scale x instead: d is η-rough if every prime factor of d is at least $(x/2)^\eta$. Such an n with prime factors $p_1 < \cdots < p_k$ has the *vector* $v(n) = (\log p_1/\log n, \dots, \log p_k/\log n)$, which lies on

$$\Delta_k(\eta) = \{u \in \mathbb{R}^k : \eta \le u_1 \le \cdots \le u_k,\ u_1 + \cdots + u_k = 1\}.$$

A *continuum function* is a family $f = (f_k)_{k \ge 1}$ of bounded functions on these simplices; we write $f(u)$ for the value at a vector u of any length, and $f(1)$ for its value at the vector (1) of a prime. The perturbations we construct have the form $w_n = F(v(n))$ on η-rough squarefree n, and $w_n = 0$ otherwise (§4).

FM translate Type I information at level γ into integral relations, their (6.2), and show that these are equivalent to a *fragmentation relation*, their (6.3) [5, Theorem 6.4]: the values of f on *free* vectors, those with every component in $[\eta, 1-\gamma)$, may be prescribed arbitrarily, and (6.3) then determines f on every other vector. We always take $\gamma = 1/2$. A vector that is not free has exactly one component $\ge 1/2$, except the vector $E = (1/2, 1/2)$, which has two. For a vector (a, β) with one large component $a \ge 1/2$, (6.3) expresses $f(a, \beta)$ through the free values on all fragmentations of a into components below $1/2$, weighted by FM's truncated Linnik kernel $L_{1/2}$:

$$f(\beta, a) = a \sum_{k \ge 2} \int_{\substack{u_1 < \cdots < u_k,\ u_1 \ge \eta \\ u_1 + \cdots + u_k = a}} L_{1/2}(u)\, f(\beta, u)\, \frac{du}{u_1 \cdots u_k},$$

with FM's projection measure. In the twisted functionals below, the factor a cancels the weight $1/a$ of the large component, so the kernels of Appendix B carry no factor a; values at vectors with a large component do. We use the following closed form of $L_{1/2}$.

Lemma 2.1 (The truncated Linnik kernel)**.** *Let $k \ge 1$, let $0 < u_i < 1/2$ for $1 \le i \le k$, let $a = u_1 + \cdots + u_k \le 1$, and let $N_{1/2}(u)$ be the number of unordered partitions of $\{1, \dots, k\}$ into two blocks each of u-sum $1/2$. Then*

$$L_{1/2}(u) = \mathbf{1}_{k=1} + (-1)^{k+1} \sum_{\substack{B \subseteq \{1,\dots,k\} \\ u_B \ge 1/2}} (-1)^{|B|} - \mathbf{1}_{a=1}\, N_{1/2}(u), \qquad u_B := \sum_{i \in B} u_i.$$

This is the case $c = 1/2$ of the generating-function identity (5.3) of [5, §5]: expand $\log V(z)$ with $V(z) = \prod_i (1 + z_i) - \sum_{u_A \ge 1/2} z^A$ and compare the coefficients of $z_1 \cdots z_k$. We give the direct count.

Proof. An unordered partition into r blocks has weight $(-1)^{r-1}(r-1)!$ in FM's (5.1), and the weights of all partitions sum to $\mathbf{1}_{k=1}$. For a prescribed block B with $u_B \ge 1/2$, summing over the partitions of its complement gives $(-1)^{k-|B|}$ (a partition containing B with r' further blocks has weight $(-1)^{r'} r'!$, and $\sum_{r'} (-1)^{r'} r'!\, S(n, r') = (-1)^n$ for the Stirling numbers $S(n, r')$, since $\sum_r (-1)^r (e^z - 1)^r = e^{-z}$); subtract these contributions. Two blocks of sum at least $1/2$ occur only

when $a = 1$, where each has sum exactly $1/2$; inclusion and exclusion adds back their weight -1 once for each unordered pair. Three such blocks are impossible. □

2.3. **The twisted functionals.** For a continuum function F and a tuple $s = (s_1, \dots, s_t)$ with $t \ge 1$, $s_i \ge \eta$ and $|s| := s_1 + \cdots + s_t < 1$, put

$$T_t(F)(s) := \sum_{k>t} (-1)^k \int_{\substack{z_{t+1}<\cdots<z_k \\ z_{t+1}+\cdots+z_k = 1-|s|}} F(s,z)\, \frac{dz}{z_{t+1}\cdots z_k},$$

$$T_0(F) := \sum_{k\ge 1} (-1)^k \int_{\substack{z_1<\cdots<z_k \\ z_1+\cdots+z_k = 1}} F(z)\, \frac{dz}{z_1\cdots z_k},$$

with FM's projection measure on the hyperplanes, and $F(s,z)$ the value at the vector with the components of s and z. The term $k = 1$ of T_0 is $-F(1)$. These are the continuum forms of the twisted sums in (H_α). For $d = p_1 \cdots p_t$ with every $p_i \ge (x/2)^\eta$, put $\lambda_i = \log p_i / \log x$ and $\lambda = \lambda_1 + \cdots + \lambda_t$. The main term of $\sum_{d|n} \mu(n) w_n$ is then

$$\int_{1-\lambda+\log(1/2)/\log x}^{1-\lambda} x^w\, q^{-1}\, T_t(F)\big(\lambda_1/q, \dots, \lambda_t/q\big)\, dw, \qquad q = \lambda + w$$

(Lemma 4.4 below), and T_0 is the case $d = 1$. The *untwisted* functionals are the same sums without the sign $(-1)^k$. For every F satisfying (6.2) at level $1/2$, the untwisted T_0 vanishes, and the untwisted $T_t(F)(s)$ vanishes for $t \ge 1$ and $|s| \le 1/2$. This is the Type I condition, and the certificates check it as a control, to the extent stated in §5.

For η-rough moduli up to x^α, the tuples that occur have $|s| = \lambda/q \le \alpha/(1 - \log 2/\log x)$, since $q \ge 1 - \log 2/\log x$, and tuples with some $s_i < \eta$ contribute nothing. So it suffices that $T_0(F) = 0$ and $T_t(F)(s) = 0$ for every $t \ge 1$ and every tuple with $s_i \ge \eta$ and $|s| \le \alpha''$, for some fixed $\alpha'' > \alpha$; the constructions of §3 give this with $\alpha'' = \alpha'$. A tuple of t components $\ge \eta$ has $|s| \ge t\eta$, so only $t \le \alpha''/\eta$ occurs.

3. The construction

3.1. **Classes and free data.** Fix $0 < \eta < \alpha' < 1/2$. A component u of a vector is *small* if $\eta \le u \le \alpha'$, *rough* if $\alpha' < u < 1/2$, and *large* if $u \ge 1/2$. Fix a partition of $(\alpha', 1/2)$ into finitely many intervals, the *bins*. A free vector u has a *class* (j, k, b): j small components, k rough components, and b the bin of its largest rough component ($b = 0$ if $k = 0$). Its *small total* is the sum τ of its small components. The free data have the form

$$f(u) = c[j,k,b] + \mathbf{1}[1 \le j \le m^*,\ \tau \le \alpha']\ w_k\, h_j(\tau), \tag{1}$$

with finitely many constants $c[j,k,b]$, a *profile* of weights w_k, and one function h_j on $[j\eta, \alpha']$ for each level $1 \le j \le m^*$. The rungs use these profiles:

rung	levels	profile
P3	$m^* = 1$	$w_k = 1$ for every k
P4	$m^* = 2$	$w_k = 1$ for every k
P5	$m^* = 3$	$w_3 = 1$, $w_k = 0$ for $k \ne 3$
P6, P7	$m^* = 7$	$w_2 = 1$, $w_3 = -1$, $w_k = 0$ otherwise

The values on all other vectors are given by FM's extension (6.3). (For P4 the level-2 function is a function of the pair's total, as written; the masked functions of §4 depend on the pair itself.)

3.2. **The census.** An η-rough squarefree modulus $d \le x^\alpha$ with $\alpha < \alpha'$ is 1 or a product of t primes, each at least $(x/2)^\eta$. Then $(x/2)^{t\eta} \le d \le x^\alpha$, so $t\eta \le \alpha/(1 - \log 2/\log x)$, which is less than α' for large x. We take $m^* = \max\{j : j\eta < \alpha'\}$. Then a modulus with $m^* + 1$ or more prime factors has $d \ge (x/2)^{(m^*+1)\eta} \ge (x/2)^{\alpha'} > x^\alpha$ for large x, so only $t \le m^*$ occurs. When $(m^* + 1)\eta = \alpha'$ exactly (P4: $3 \cdot \frac{1}{18} = \frac{1}{6}$; P5: $4 \cdot \frac{1}{20} = \frac{1}{5}$; P7: $8 \cdot \frac{1}{30} = \frac{4}{15}$), a function h_{m^*+1} could live only at the single point $\tau = \alpha'$, and we set it to zero: the *null level*. The census then has no slack and uses $\alpha < \alpha'$ strictly.

3.3. **The rows and the forcing.** Unmasked, the free data (1) see the small part of a vector only through j and τ. So the twisted functional at a tuple of j small components with total σ depends only on σ, and it has the form

$$T_j(\sigma) = A_j(\sigma) \cdot c + B_j(\sigma)\, h_j(\sigma) + \sum_{i \ge 1,\ j+i \le m^*} C_{j,i}[h_{j+i}](\sigma), \qquad \sigma \in [j\eta, \alpha'], \tag{2}$$

where $A_j(\sigma) \cdot c$ collects the constants, the *diagonal* $B_j(\sigma)$ collects the completions of the tuple that add no small component, weighted by w_k, and $C_{j,i}$ collects those that add exactly i small components. No h_J with $J < j$ enters, since the row's vectors have at least j small components. If B_j does not vanish on $[j\eta, \alpha']$, then

$$h_j := -\big(A_j \cdot c + \textstyle\sum_i C_{j,i}[h_{j+i}]\big)/B_j, \qquad j = m^*, m^* - 1, \dots, 1,$$

defines the functions top-down and makes every row $T_j \equiv 0$ hold *exactly*; each h_j is linear in c. This is the *forcing*. What remains to choose are the finitely many constants. The kernels and the rows in full are in Appendix B.

For the profiles above the diagonal has a sign. For P5, B_j uses only completions with exactly three rough parts, and the kernels of Appendix B give $B_j(\sigma) = 2(-1)^{j+1}(R_1 * R_2\mathbf{1}_{[1/2,1]})(1 - \sigma)$, which is nonzero exactly when $1 - \sigma > 1/2 + \alpha'$, hence on all of $[j\eta, \alpha']$ when $\alpha' < 1/4$; at $\alpha' = 1/5$ its minimum is at $\sigma = \alpha'$ (value $0.8250801\ldots$), and the certificate's computation, as recorded by the audit's replay of it (§10), gives $\min_j |B_j| \ge 0.825079$. For the weighted profile there is also a closed form.

Lemma 3.1 (The weighted diagonal)**.** *Let $\alpha' \ge 1/4$, and let $R_k = r^{*k}/k!$, where $r(u) = u^{-1}$ on $(\alpha', 1/2)$ and $r(u) = 0$ elsewhere. For the profile $w_2 = 1$, $w_3 = -1$,*

$$B_j(\sigma) = (-1)^j\big[2R_2(t) + 6R_3(t)\big], \qquad t = 1 - \sigma.$$

In particular $(-1)^j B_j \ge 0$. If $\alpha' < 1/3$, then B_j has the sign $(-1)^j$ and does not vanish on $[j\eta, \alpha']$; if $\alpha' \ge 1/3$, then $B_j \equiv 0$ on the nonempty interval $[\max(j\eta, 1 - 2\alpha'), \alpha']$.

Proof. A completion of the tuple that adds no small component either is free, with k rough parts, or has one large part that fragments into rough parts. Rough parts exceed $\alpha' \ge 1/4$, so R_2 and R_3 vanish on $[0, 1/2)$, and $R_1\mathbf{1}_{[1/2,\infty)} = 0$. Hence the Lemma 2.1 kernels of a large part fragmenting into k_y rough parts are $\Lambda_2^0 = -R_2$ and $\Lambda_3^0 = 2R_3$. With the row sign $(-1)^{\dim}$, the contributions are, for $k = 2$, $(-1)^j R_2 + (-1)^{1+j}\Lambda_2^0 = 2(-1)^j R_2$, and for $k = 3$, $(-1)^{j+1}R_3 + (-1)^{1+j}\Lambda_3^0 + (-1)^j(R_1 * \Lambda_2^0) = (-1)^{j+1}(1 + 2 + 3)R_3$, using $R_1 * R_2 = 3R_3$. Weighting by $w_2 = 1$ and $w_3 = -1$ gives the claim. Since $R_2(t) > 0$ for $t \in (2\alpha', 1)$, which contains $1 - \sigma$ for every $\sigma \le \alpha'$ when $\alpha' < 1/3$, the diagonal then does not vanish. If $\alpha' \ge 1/3$, then $R_3(t) = 0$ for $t < 1 \le 3\alpha'$, and $R_2(t) = 0$ for $t \le 2\alpha'$, that is, for $\sigma \ge 1 - 2\alpha'$; for $\alpha' \ge 1/3$ these are the rows of Lemma 8.1. □

3.4. **The parity solve and the target inequalities.** The row T_0 is linear in c, with the h_j entering through c. Choose a direction g supported on classes with $j = 0$ (no small component). Then $A_j \cdot g = 0$ for every $j \ge 1$, since rows with a small prime meet only classes with $j \ge 1$, so moving along $c + tg$ leaves every h_j unchanged, and the solve $T_0(c + tg) = 0$ is affine in t. The directions used are listed in Table 2: in the cases with $m^* \le 3$, g is the unit vector of one class with

no small and three rough components, and $g = e_{(0,3,2)} + e_{(0,3,3)}$ in the two cases with the weighted profile.

With $c^* = c + t^* g$, the certificate (§5) proves:

$$f(1) < 0 \qquad \text{and} \qquad f(u) \geq (1-m)\, f(1) \ \text{ for every non-prime } \eta\text{-rough vector } u \neq E, \tag{3}$$

for an explicit $m > 0$, including the vectors with one large component, whose values come from (6.3). In the two cases with $m^* = 1$ the certified statement is narrower: it covers every such vector with no subsum equal to $1/2$. No lower bound is asserted at E, nor in those two cases on the other vectors with a subsum equal to $1/2$; the mask of §4 makes the masked extension vanish there; Theorem 4.1 uses (3) only at vectors with no subsum equal to $1/2$.

3.5. **The constants.** The constants are found by a linear programme: minimise $\sum |c|$ subject to an upper bound on $f(1)$, to the positivity rows at grid points, to $T_0 = 0$ and to box constraints, on an exact oracle at a moderate mesh. This is provenance only. The certificate takes the constants as exact rationals, the binary64 values of the stored solution, and does not use the programme.

4. From a certified continuum function to sequences

This section proves that the continuum statement of hypothesis (2) below, together with the non-vanishing of the diagonals and of $T_0(g)$, gives the theorem. The argument is FM's construction procedure, the proof of [5, Theorem 6.3], with their mask, extended by a re-solve of the forced functions after masking and by a reduction of the twisted sums to the rows of §3.

Theorem 4.1. *Let $0 < \eta < \alpha' < 1/3$ with $\eta \leq 1/4$, let $m^* = \max\{j : j\eta < \alpha'\}$, and let free data of the form (1) be given, with the h_j forced as in §3 and the constants $c^* = c + t^* g$ solving $T_0 = 0$. Suppose:*

1. *$|B_j(\sigma)| \geq \beta_0 > 0$ for $1 \leq j \leq m^*$ and $\sigma \in [j\eta, \alpha']$, and $|T_0(g)| \geq \tau_0 > 0$;*
2. *$f(1) < 0$, and for some $m > 0$ the inequality $f(u) \geq (1-m)f(1)$ of (3) holds at every non-prime η-rough vector u with no subsum equal to $1/2$ (such u differ from E); and $|f| \leq C$ on all vectors.*

Then for every $0 < \nu < \eta$ and every $0 \leq \alpha < \alpha'$ there are $\varepsilon > 0$ and a bounded function F_ε on η-rough vectors, both independent of B and x, with $F_\varepsilon(1) = -1$ and $F_\varepsilon \geq -1$, such that the following holds. For $x \geq 2$ let $a_n = a_n^{(x)} := 1 + F_\varepsilon(v(n))$ for squarefree $n \in (x/2, x]$ whose prime factors are all at least n^η, $a_n := 1$ for the other $n \in (x/2, x]$ and $a_n := 0$ for $n \notin (x/2, x]$, and let $b_n := \mathbf{1}_{(x/2,x]}(n)$. Then $0 \leq a_n \leq 1 + \sup|F_\varepsilon|$ for every x, and $a_p = 0$ for every prime $p \in (x/2, x]$. For every $B > 0$ there is x_0, depending on B and on the data above, such that for every $x \geq x_0$ the sequence $w_n = a_n - b_n$ satisfies (I) at $\gamma = 1/2$, (II) on $1 < \ell \leq x^\nu$, (w) and (H_α). In particular $C^-(1/2, 0, \nu \mid \mathrm{H}_\alpha) = 0$.

4.1. **Setup and the mask.** Fix $\alpha < \alpha'$ and $0 < \nu < \eta$. Fix $\varepsilon > 0$ with $\nu + \varepsilon < \eta$ and $\varepsilon < 1/2 - \alpha'$, and so small that the $O(\varepsilon)$ bounds of Steps 1–3 give $|B_{j,\varepsilon}| \geq \beta_0/2$ at every level, $|T_{0,\varepsilon}(g)| \geq \tau_0/2$, $f_\varepsilon(1) < 0$ and the margin inequality of Step 3. These thresholds depend only on η, α', the free data, g, C, β_0, τ_0, m and $f(1)$, not on x; ε is fixed before $x \to \infty$. Put $I_\varepsilon = [\nu, \nu + \varepsilon] \cup [1/2 - \varepsilon, 1/2 + \varepsilon]$ and $\chi_\varepsilon(u) = \mathbf{1}[u \text{ has no subsum in } I_\varepsilon]$. In Lemma 4.2 implied constants depend only on η and $\sup|g_0|$ [5, (6.6)]; in Steps 1–3 they depend on η, α', C, the free data (c^*, the weights w_k, the h_j), g, β_0 and τ_0, but not on ε, x or the tuple; in Step 4 and Lemma 4.4 they also depend on ε, through the bound and the Lipschitz data of F_ε.

Lemma 4.2 (Masking costs $O(\varepsilon)$, uniformly)**.** *Let g_0 be bounded free data, piecewise Lipschitz on finitely many polytopes, let g be its extension by (6.3), and let g_ε be the extension of $g_0\chi_\varepsilon$. Then:*

(a) *g_ε satisfies (6.2), the continuum Type I relation at level $1/2$ [5, Theorem 6.4];*
(b) *$g_\varepsilon(u) = 0$ whenever u has a subsum in I_ε, and $|g(u) - g_\varepsilon(u)| \ll \varepsilon$ otherwise [5, (6.6)];*

(c) *for every row functional R (the functional T_0, and the twisted and untwisted functionals at every tuple s of $t \le m^*$ small components with $|s| \le \alpha'$), $|R(g) - R(g_\varepsilon)| \ll \varepsilon$, uniformly in the tuple.*

Proof. (a) is [5, Theorem 6.4]; (b) is [5, (6.6)] and the remark preceding it in [5, §6.2]. For (c), each row is a finite sum over η-rough ordered completions z of the tuple s to total 1, with uniformly bounded densities and total masses. The half $[\nu, \nu+\varepsilon]$ of I_ε is vacuous, since the subsums of an η-rough vector are 0 or at least $\eta > \nu + \varepsilon$. A subsum that contains no completion coordinate is at most $|s| \le \alpha'$, and one that contains all of them is at least $1-|s| \ge 1-\alpha'$. Since $\alpha' < 1/2-\varepsilon$, neither meets $[1/2-\varepsilon, 1/2+\varepsilon]$. Every other subsum contains some, but not all, completion coordinates. Holding the other coordinates fixed, vary one included and one excluded completion coordinate with their sum fixed: the subsum moves with slope 1, so the set of such variations on which it lies in I_ε has length $O(\varepsilon)$. Summing over the finitely many index sets, the set U_ε of completions with a subsum in I_ε has measure $O(\varepsilon)$, uniformly in the tuple. For completions with one component, U_ε is empty; this includes the prime vector in T_0. On U_ε the masked extension vanishes and the unmasked one is bounded, and off U_ε part (b) applies. This proves (c). □

4.2. **Step 1: the masked forced functions.** Unmasked, the rows see the small part only through (j, τ). The mask sees subsums such as $s_i + r$ separately, so the masked rows depend on the tuple. FM allow free data that are bounded and Lipschitz on finitely many polytopes (Lemma 4.3), so we let the forced functions depend on the tuple:

$$f_{0,\varepsilon}(u) := c^*[j,k,b]\,\chi_\varepsilon(u) + \mathbf{1}\big[1 \le j \le m^*,\ s_1 + \cdots + s_j \le \alpha'\big]\, w_k\, h_{j,\varepsilon}(s_1, \dots, s_j)\,\chi_\varepsilon(u),$$

where $s_1, \dots, s_j$ are the small components of u. The masked coefficients are the actual tuple rows of the (6.3) extensions of the masked basis data, with the mask kept inside every fragmentation integral. They are invariant under permutations of the tuple, so $h_{j,\varepsilon}$ is symmetric and the free data are well defined on multisets.

Solve for $j = m^*, m^*-1, \dots, 1$ in turn. In the masked row $T_{j,\varepsilon}(s_1, \dots, s_j)$, the function $h_{j,\varepsilon}$ enters only through free data whose small part is exactly the tuple, and whose rough part, counted after fully fragmenting any large component, has k parts with $w_k \ne 0$; its coefficient is the masked diagonal $B_{j,\varepsilon}(s) := \sum_k w_k B^{(k)}_{j,\varepsilon}(s)$, where $B^{(k)}_{j,\varepsilon}(s)$ is the contribution of the completions whose rough part has k parts. The functions $h_{J,\varepsilon}$ with $J > j$ enter through completions adding exactly $J - j$ small components, with the total of all J small components at most α' (as free components, inside the remaining free part, or as fragments of the large part), and they are already defined. The functions with $J < j$ do not enter. A vector with $m^* + 1$ or more small components carries no h. So

$$T_{j,\varepsilon} = A_{j,\varepsilon} \cdot c^* + B_{j,\varepsilon}\, h_{j,\varepsilon} + W^{(j)}_\varepsilon[h_{j+1,\varepsilon}, \dots, h_{m^*,\varepsilon}].$$

By Lemma 4.2(c) and induction on the solved levels, $|B_{j,\varepsilon}(s) - B_j(|s|)| \ll \varepsilon$, $|A_{j,\varepsilon} - A_j| \ll \varepsilon$ and $|W^{(j)}_\varepsilon[h_{j+1,\varepsilon}, \dots, h_{m^*,\varepsilon}](s) - W^{(j)}[h_{j+1}, \dots, h_{m^*}](|s|)| \ll \varepsilon$, uniformly in the tuple. Since $|B_j| \ge \beta_0$, for small ε

$$h_{j,\varepsilon} := -\big(A_{j,\varepsilon} \cdot c^* + W^{(j)}_\varepsilon\big)/B_{j,\varepsilon}$$

is well defined, $|h_{j,\varepsilon}(s) - h_j(|s|)| \ll \varepsilon$ uniformly, and $T_{j,\varepsilon} \equiv 0$ at every j-tuple with total $\le \alpha'$. Because the masked data are defined with the same weights w_k, each masked row has a single unknown, whose coefficient is a fixed combination of masked diagonals; no consistency condition between the parts with different k arises.

4.3. **Step 2: parity.** The direction g has no small component, so $A_j \cdot g = 0$ for $j \ge 1$, masked or not, and moving along $c^* + tg$ leaves every $T_{j,\varepsilon}$ unchanged. Now $T_{0,\varepsilon}(c^*) = O(\varepsilon)$ and $T_{0,\varepsilon}(g) = T_0(g) + O(\varepsilon)$ with $|T_0(g)| \ge \tau_0$. So $t_\varepsilon := -T_{0,\varepsilon}(c^*)/T_{0,\varepsilon}(g) = O(\varepsilon)$, and $c_\varepsilon := c^* + t_\varepsilon g$ gives $T_{0,\varepsilon} = 0$ and $T_{j,\varepsilon} \equiv 0$ for every $1 \le j \le m^*$. Let f_ε be the resulting masked extension.

4.4. **Step 3: positivity and normalisation.** On vectors with a subsum in I_ε the masked extension is zero. This includes E, since every fragmentation of E has a subsum equal to $1/2$, and E is the only vector with two components $\geq 1/2$. On every other vector, the boundedness of (6.3), Lemma 4.2(b) and the $O(\varepsilon)$ bounds above give $f_\varepsilon = f + O(\varepsilon)$ uniformly. The prime vector (1) is unmasked, so $f_\varepsilon(1) = f(1) + O(\varepsilon) < 0$, and every unmasked non-prime vector has no subsum in $[1/2 - \varepsilon, 1/2 + \varepsilon]$, in particular none equal to $1/2$, so it differs from E and has the margin of hypothesis (2). We may assume $m \leq 1$, since (3) for m implies it for every smaller $m > 0$. So for small ε, $f_\varepsilon(u) \geq (1 - m/2) f_\varepsilon(1)$ on every non-prime vector: on the unmasked ones by that margin, and on the masked ones because there $f_\varepsilon = 0$ while $(1 - m/2) f_\varepsilon(1) < 0$. Put $F_\varepsilon := f_\varepsilon / |f_\varepsilon(1)|$. Then $F_\varepsilon(1) = -1$, $F_\varepsilon \geq -1 + m/2$ on every other vector, $|F_\varepsilon| \leq C'$, $T_0(F_\varepsilon) = 0$, and $T_t(F_\varepsilon)(s) = 0$ at every tuple of $t \leq m^*$ small components with $|s| \leq \alpha'$; and F_ε satisfies (6.2) by Lemma 4.2(a), whose appeal to [5, Theorem 6.4] uses axiom (b) of [5, Definition 6.2], supplied by Lemma 4.3.

Lemma 4.3 (Regularity)**.** *F_ε is symmetric and satisfies FM's Definition 6.2(a, b): in each dimension it is a finite sum of functions supported and Lipschitz on convex polytopes.*

Proof. Dimensions are bounded by $1/\eta$. The class, bin and profile conditions (the number of rough parts, the bin of the largest one), the ordering conditions, the mask and the partition inequalities defining $L_{1/2}$ are finitely many affine conditions. For each level in turn, expand the finite fragmentation and Lemma 2.1 subset sums that define $A_{j,\varepsilon}$, $B_{j,\varepsilon}$ and $W_\varepsilon^{(j)}$, keeping χ_ε. The domains are polytopes with fixed normals and affine dependence on the tuple, every coordinate that appears in a denominator is at least η, and the integrals of bounded integrands over polytopes with fixed normals moving affinely with the tuple are Lipschitz, with chambers created only by point evaluations and by the affine walls of the integrand. $B_{j,\varepsilon}$ is bounded away from zero, so $h_{j,\varepsilon}$ is piecewise Lipschitz; refine the chambers before the next level. The extension (6.3) and the normalisation preserve this. □

4.5. **Step 4: the sequences.** Put $w_n := F_\varepsilon(v(n))$ for squarefree $n \in (x/2, x]$ whose prime factors are all at least n^η and $w_n := 0$ for the other $n \in (x/2, x]$, and let $a_n := 1 + w_n \geq 0$ and $b_n := 1$ on $(x/2, x]$, with $a_n = b_n = 0$ elsewhere, as in Theorem 4.1. Then $a_p = 1 + F_\varepsilon(1) = 0$ for every prime $p \in (x/2, x]$. (I) at $\gamma = 1/2$ follows from the argument for [5, (6.7)–(6.9)] in the proof of [5, Theorem 6.3], applied to the already masked F_ε (Lemma 4.3); the strip $[1/2 - \varepsilon, 1/2 + \varepsilon]$ supplies their case $m > x^{\gamma - \varepsilon}$, and no second mask is applied. (II) on $1 < \ell \leq x^\nu$ holds because every perturbation term vanishes: such an ℓ has a prime factor at most x^ν, and $x^\nu < (x/2)^\eta$ for large x, so ℓ divides no element of the support of w. (w) holds since a_n is bounded uniformly in x.

For (H_α), which is imposed on $w = a - b$, we bound $\sum_{d|n} \mu(n) w_n$ for $d \leq x^\alpha$ (by §2, Hooley's form for a_n then follows as well). A non-squarefree d contributes 0. If d has a prime factor below $(x/2)^\eta$, the sum is 0. If d has $m^* + 1$ or more prime factors $\geq (x/2)^\eta$, then $d \geq (x/2)^{(m^*+1)\eta} > x^\alpha$ for large x, since $(m^* + 1)\eta \geq \alpha' > \alpha$. For $d = 1$, the proof of Lemma 4.4 with $t = 0$ ($\lambda = 0$, $q = w$ and $\int_{1 - \log 2/\log x}^1 x^w \, dw/w = \mathrm{Li}(x) - \mathrm{Li}(x/2)$) gives $\sum \mu(n) w_n = (\mathrm{Li}(x) - \mathrm{Li}(x/2))\, T_0(F_\varepsilon) + O(x e^{-(\log x)^{1/5}}) = O(x e^{-(\log x)^{1/5}})$. For $d = p_1 \cdots p_t$ with $1 \leq t \leq m^*$, every $p_i \geq (x/2)^\eta$ and $d \leq x^\alpha$, Lemma 4.4 gives $\ll (x/d) e^{-(\log x)^{1/5}}$. Summing,

$$\sum_{d \leq x^\alpha} \Big| \sum_{d|n} \mu(n) w_n \Big| \ll x e^{-(\log x)^{1/5}} \sum_{t=0}^{m^*} \Big(\sum_{p \leq x} \frac{1}{p} \Big)^t \ll x e^{-(\log x)^{1/5}} (\log\log x)^{m^*} \ll_B x (\log x)^{-B}.$$

Lemma 4.4 (Twisted sums at products of t primes)**.** *Let $1 \leq t \leq m^*$. Uniformly over primes $p_1 < \cdots < p_t$ with $p_1 \geq (x/2)^\eta$ and $d = p_1 \cdots p_t \leq x^\alpha$,*

$$S_d^\mu := \sum_{\substack{x/2 < n \leq x \\ d|n}} \mu(n) w_n \ll_\varepsilon \frac{x}{d} \, e^{-(\log x)^{1/5}}.$$

Proof. We have $S_d^\mu = \mu(d)\sum_r \mu(r)w_{dr} = (-1)^t \sum_r \mu(r)w_{dr}$, over squarefree r coprime to d with $x/(2d) < r \le x/d$. This is FM's sum $S_m(c,d')$ of [5, (6.8)] with their m equal to our d, $\lambda_i = \log p_i/\log x$, $\lambda = \sum \lambda_i \le \alpha < 1/2 - \varepsilon$ and $(c,d') = (1/2,1)$, except for the factor $(-1)^t\mu(r)$; FM's (6.8) holds for every $m \le x^{1/2-\varepsilon}$, whatever its number of prime factors. FM's pieces fix $k = \omega(dr)$, the polytope piece of F_ε and the meshing; in each piece $\omega(r) = k - t$, so $\mu(r) = (-1)^{k-t}$. FM's Lemma 5.11, applied to each piece (normalised by a constant $K_\varepsilon \ge 1$ controlling the sup norm and the Lipschitz constant, then multiplied back; its hypotheses hold, since $\eta \le 1/4$ and d is squarefree with $d \le x^\alpha < x^{1-\eta}$), gives

$$S_d^\mu = \sum_{k \ge t+1} (-1)^k \,\mathrm{main}_k + O_\varepsilon\big((x/d)\, e^{-(\log x)^{1/5}}\big),$$

where main_k is FM's untwisted main term in total dimension k and the sign includes $\mu(d) = (-1)^t$. Their sum is the integral

$$\int_{1-\lambda+\log(1/2)/\log x}^{1-\lambda} x^w \sum_{k \ge t+1} (-1)^k \int_{\substack{u_{t+1}+\cdots+u_k = w \\ u_{t+1} < \cdots < u_k}} \frac{F_\varepsilon\big(\lambda_1/q, \ldots, \lambda_t/q, u_{t+1}/q, \ldots, u_k/q\big)}{u_{t+1}\cdots u_k}\, d\boldsymbol{u}\, dw,$$

with $q = \lambda + w$. Rescale $u_i = qz_i$: the projection measure contributes q^{k-t-1} and the denominator $q^{-(k-t)}$, so the inner sum is exactly $q^{-1}\,T_t(F_\varepsilon)(\lambda_1/q, \ldots, \lambda_t/q)$, with no factorial or logarithmic factor. If some $\lambda_i/q < \eta$, the integrand vanishes, since F_ε is supported on η-rough vectors. Otherwise $\sum_i \lambda_i/q = \lambda/(\lambda + w) \le \alpha(1 + o(1)) < \alpha'$ for large x, and $T_t(F_\varepsilon) = 0$ at that tuple by Step 3. Uniformity in $(p_1, \ldots, p_t)$ is that of Lemma 5.11. □

This proves Theorem 4.1. □

5. The certificates

5.1. **What a certificate establishes.** Fix a case: the parameters (η, α'), the profile weights w_k, the bins, and the class constants c, which are binary64 numbers used at their exact values. The forced functions h_j and the parity correction are then determined (§3). A certificate is a computation in exact integer arithmetic, in integer arithmetic with directed rounding, or, in the separate implementation of §5.2, in binary64 interval arithmetic with exact integer convolutions, that encloses every quantity below and checks the following strict inequalities.

- (C1) The weighted diagonals do not vanish: $|B_j(\sigma)| \ge \beta_0 > 0$ for $1 \le j \le m^*$ on the forcing range, with the signs stated in Lemma 3.1 where it applies.
- (C2) The parity direction is effective: $|T_0(g)| \ge \tau_0 > 0$. The solution t^* of $T_0(c + tg) = 0$ is then enclosed in an interval.
- (C3) Uniformly over t in that interval, $f(1) \le U < 0$, and $f(u) \ge L$ for every non-prime η-rough vector $u \ne E$, free or with one component of size at least $1/2$; in the two cases with $m^* = 1$, for every such vector with no subsum equal to $1/2$, as stated after (3). On a vector with a large component a and remaining part β, the value bounded is $f(\beta, a)$ of §2, including its factor a.

The *margin* is $1 - L/U$, and the lift (Theorem 4.1) is applied with any fixed m below it. Indeed, if $0 < m \le 1 - L/U$, then $f(u) \ge L = (1 - (1 - L/U))\,U \ge (1 - m)\,U \ge (1 - m)f(1)$, since $U < 0$, $U \ge f(1)$ and $m \le 1 - L/U < 1$, every certified L being negative. The second-order engine and the exact certificate of $(1/18, 1/6)$ also enclose, for the function of record, the parts of each untwisted row T_j, $j \ge 1$ (its constant part, its diagonal coefficient, and each coupling applied to the forced h_{j+i}), and the untwisted T_0. All of these vanish identically by construction, and each enclosure is checked to contain 0. The engine makes this part of its verdict; the exact certificate records it. The exact certificate of $(1/11, 1/10)$ and $(1/14, 1/8)$ does the same for the constant part and the h_1-coefficient of the untwisted row T_1, but does not compute the untwisted T_0. The separate

implementation checks the untwisted rows only in a separate mode, on random data; its certificate runs do not include that mode, and the outputs of that mode are not in the supplement. These checks test the machinery and are not used in the proofs.

The enclosed quantities are built from the kernels $1/u$ on $[\eta, \alpha']$ and on $(\alpha', 1/2)$ by convolution, restriction to bins, Lemma 2.1's fragmentation formula (over all subsets, those of small fragments alone included), and the pointwise forcing $h_j = -(\cdots)/B_j$. Most of them are functions of one variable, the total size of a group of components. Two further variables are handled in blocks: in the couplings and Ω terms, the subset sum ς of Appendix B, the total of those small fragments of a large component that belong to the sub-collection B of Lemma 2.1 (the rough members of B must then total at least $1/2 - \varsigma$; in this section and in Lemma 5.6 this variable is written τ and called the threshold); and, in the positivity bounds for vectors with a large component, the small total σ_β of the part outside it. A certificate discretises the variable on the grid $\{n/M\}$ and carries rigorous upper and lower enclosures through every operation; §5.4 states the enclosure lemmas they rest on. Every enclosure is outward. Apart from terms bounded over blocks of fixed width, the first-order methods lose $O(1/M)$ in width and the second-order engine $O(1/M^2)$. The block terms lose an amount of first order in the block width, which does not shrink with M; which terms they are is stated for each implementation below.

5.2. **First-order certificates.** Each kernel is enclosed between first-order bound functions on the grid (step or piecewise linear), using the monotonicity of $1/u$. Their convolutions are exact convolutions of integer vectors, and quotients and products are rounded outward after every operation. Certificates of this kind were written from the mathematical text of each case, not from the code that produced its constants. They certify $(1/14, 1/8)$ and $(1/11, 1/10)$ (the certificates of record for those cases) and $(1/18, 1/6)$.

In the certificate of $(1/18, 1/6)$, the couplings and Ω terms of h_2 (in T_1, T_0, $f(1)$ and the extended vectors) are bounded at every M over 50 threshold blocks of width $(\alpha' - \eta)/50$, by their values at the block ends; the positivity bounds for extended vectors whose β-part has one small and at least one rough component range over 50 blocks of σ_β; and free vectors with two small components are bounded below by the smaller of c and c plus the minimum of h_2 over $[2\eta, \alpha']$. Its margins 0.0067, 0.0179, 0.0290 at $M = 108{,}000$, 144,000, 216,000 fit $0.0515 - 4835/M$ to within $6 \cdot 10^{-5}$, so they extrapolate to about 0.051; the separate implementation's margins extrapolate to about 0.062. That implementation encloses T_0 and $f(1)$ without blocks. Fixed groups of σ_β and prefix ranges of h_1 and h_2 enter only its positivity bounds for extended vectors whose β-part has one or two small components and at least one rough one, and these do not bind.

A separate implementation (§5.6), written from the mathematical text without sight of the code of these certificates or of the second-order engine, certifies $(1/18, 1/6)$ and $(1/20, 1/5)$ again. It is a first-order certificate in validated binary64 arithmetic: every convolution is an exact integer product of integer images of its operands, scaled by exact powers of two and rounded outward, and every other operation is done in binary64 interval arithmetic, with the error control of Lemma 5.8. The step $1/M$ enters only through division by the exact integer M. No library transform and no transcendental function enters any enclosure; in the $(1/20, 1/5)$ code a library $\log_2$ is used only to choose indices in a range-extremum table, and its choices were checked against exact integer bit lengths for every length that occurs. This is what separates the implementation from the float computations of §5.6, whose convolutions are FFTs bounded a priori. In its $(1/20, 1/5)$ run at $M = 432{,}000$, the extended vectors whose part β has one or two small and one or two rough components are bounded on the coarser grid $\{n/72{,}000\}$: the forced functions enter through unions of their enclosures over the six fine cells of each coarse cell, and σ_β ranges over 16 blocks. These lower bounds are valid but wider. Those classes carry no term in the parity direction. The least margin, 0.01609, is set at the upper end of the parity segment by a class computed on the full grid; one of the coarse classes, $(2, 1, 1)$, gives the least value at its lower end. The separate

implementation does not carry t as an interval through (C3). It evaluates $f(1)$ and the lower bounds of $f(u)$ at the two ends t_-, t_+ of its enclosure of t^*, obtains $U_\pm < 0$ and $L_\pm$ there, and reports $\min(1 - L_-/U_-,\ 1 - L_+/U_+)$. This suffices. Since g leaves every h_j unchanged (§3), $f(1)$ and every $f(u)$ are affine in t. Write $t^* = (1-\theta)t_- + \theta t_+$ with $\theta \in [0,1]$, and let $0 < m$ be at most the reported value. Every certified $L_\pm$ is negative, so $m < 1$ and $L_\pm \geq (1-m)U_\pm$. Hence at t^*, $f(u) \geq (1-\theta)L_- + \theta L_+ \geq (1-m)\big((1-\theta)U_- + \theta U_+\big) \geq (1-m)f(1)$, and $f(1) < 0$ there. In the three cited runs ($M = 144{,}000$, $288{,}000$ and $432{,}000$) the smaller L and the larger U occur at the same end, t_+, so the reported value also equals $1 - L/U$ for the uniform bounds $L = \min L_\pm$, $U = \max U_\pm$ of (C3).

5.3. **The second-order engine.** To reach the margins at $(1/20, 1/5)$ and beyond at feasible resolution, a second-order engine was built.

- **Representation.** Each enclosure is a pair of functions that are affine on every grid cell and may jump at grid points. They are stored as integers scaled by 10^{18}.
- **Convolution.** The convolution of two such functions is a continuous piecewise cubic. Its values at the grid points are exact integer convolutions, three FLINT products per pair. On each cell it differs from its chord by at most $1/(8M^2)$ times a bound for its second derivative, which is itself a convolution of the derivatives, jumps included, and is enclosed exactly (Lemma 5.2). Signs are handled by splitting each operand into positive and negative parts. On the cells where an operand changes sign the loss is $O(1/M)$ on one cell, so $O(1/M^2)$ in any integral.
- **Products and quotients** on a cell are enclosed through the four corner products and a second-derivative slack. For the forcing quotient $-N/B$ with B of fixed sign on the cell, the slack is computed exactly in rationals (Lemma 5.4).
- **First-order terms that remain.** For each number $b \geq 1$ of small fragments in the subcollection, the threshold $\tau \in [b\eta, \alpha']$ ranges over at most NB blocks of $\lceil(\alpha' - b\eta)M/NB\rceil$ cells; the last block may be shorter. In the couplings and Ω terms (the contributions of the forced functions inside a fragmented large component) whose chains carry no rough factor, the blocks are bounded by their narrow and wide ends; elsewhere the block start is exact and the remainder is bounded by the window ranges of Lemma 5.6(b). In the positivity bounds of extended vectors, σ_β ranges over blocks: at most NBQ of them for $j_\beta \geq 3$ and $NBQLO$ for $j_\beta \leq 2$, or $NBQJ$ set per level; in the runs of `cert2_v2.py`, which has no separate counts, at most NB (40 for $(1/20, 1/5)$ and 10 for the second-order run of $(1/18, 1/6)$). Their widths are fixed by the block counts, not by M. The positivity blocks are local. The coupling blocks enter the forced h_j and propagate down the forcing chain; for the cases certified here they are negligible. The positivity loop over σ_β uses range queries over precomputed extremes of h_j (a fast path whose outward enclosure was audited).

The engine evolved in steps. Each step was audited before a run on it was cited.

(1) *Weighted profiles.* h_j enters the free data on vectors with k rough components with weight w_k, and the diagonal is the weighted one of Lemma 3.1. Without weights the code reproduces the earlier unweighted runs exactly; with all weights equal to 1 on their support the weighted formulas reduce to the unweighted ones.

(2) *Adjoint enclosures.* T_0 and $f(1)$ are linear functionals of the forced functions. Instead of integrating the interval widths of the forcing chain, which grow by a bounded factor per level, the engine also solves the adjoint recursion of Lemma 5.5 and integrates each width once against the adjoint solution. Both enclosures are valid; the engine checks that they intersect and uses the intersection. The $(1/30, 1/4)$ run uses the adjoint for T_0 only; the $(1/30, 4/15)$ run uses it for T_0 and $f(1)$.

(3) *Adaptive blocks.* At $(1/30, 4/15)$ one extended class had a floor independent of M, from the block over which the β-part's total σ_β ranges (it enters h_2 and the shifted h_{2+i} in the Ω

terms). Finer blocks for the affected levels, and an adaptive partition with explicit targets, reduce the floor below what the margin needs. In the cited run the σ_β-range of that class is cut into $NBQLO = 64$ base blocks, 54 of which meet the class's feasible range, and 7 of these were split into 8. The targets affect work and enclosure width, never the pass criterion.

(4) *Block-by-block evaluation.* The extended classes with σ_β blocks are evaluated one block at a time, holding one block's Ω terms in memory instead of all of a class's blocks. For the same partition this gives the same bounds (only the reported minimiser can differ in exact ties), and it is what lets fine and adaptive partitions fit in memory.

Steps 3 and 4, with the per-level block counts $NBQJ$, were audited together, as one engine step, before the $(1/30, 4/15)$ run was cited. From the start, the engine bounds each free class (j, k, b) with $1 \le j \le m^*$ and $w_k \ne 0$ below by $c + w_k$ times the extreme of h_j (its minimum if $w_k > 0$, its maximum if $w_k < 0$) over the whole closed interval $[j\eta, \alpha']$, not over the class's realisable range. For $j \ge 2$ it also takes c itself as a candidate, the value on the vectors of the class with $\tau > \alpha'$, which carry no h_j, and keeps the smaller of the two; it does this on every such class, whether or not the class has vectors on both sides of α'. For $j = 1$ no such candidate is needed, since a single small component never exceeds α'. The other free classes ($j = 0$, $j > m^*$ or $w_k = 0$) are bounded by c. This is a valid, conservative enclosure: the margin stays valid (it can only be smaller than a bound over realisable ranges would give), but the class reported as binding may be a label that no vector realises.

5.4. **The enclosure lemmas.** Let $x_n = n/M$ and $K_n = [x_n, x_{n+1}]$ for $0 \le n < M$. A function on $[0, 1)$ is *cellwise affine* if it is affine on each open cell; it may jump at grid points, and it is taken to vanish outside $[0, 1)$. An *enclosure* of f is a pair $F^- \le F^+$ of cellwise affine functions with $F^- \le f \le F^+$ on every open cell. The exact-integer and second-order implementations store their end values as integers over 10^{18}, rounded down for lower bounds and up for upper bounds. The separate implementation stores binary64 interval endpoints and uses outward integer images for its exact convolutions. Step and continuous piecewise-linear bounds are special cases. The certificates combine exact linear identities, restrictions, reflections, integrals and range hulls with the enclosure operations below. The proofs are in Appendix C.

Lemma 5.1 (Kernels)**.** *For $1 \le i < M$ and $u \in K_i$, $M/(i+1) \le 1/u \le M/i$ and $\ell_i(u) - M/(4i^3) \le 1/u \le \ell_i(u)$, where ℓ_i is the chord of $1/u$ on K_i.*

Lemma 5.2 (Convolutions and correlations)**.** *Let F and G be cellwise affine, with end values $a_n = F(x_n+)$, $b_n = F(x_{n+1}-)$ and c_m, d_m for G, increments $\Delta a_n = b_n - a_n$, $\Delta c_m = d_m - c_m$, and jumps $J^F_n = a_n - b_{n-1}$, $J^G_m = c_m - d_{m-1}$ (with $b_{-1} = d_{-1} = 0$). Then $s = F * G$ is continuous, and:*

(a) $s(x_k) = \frac{1}{6M} \sum_{n+m=k-1} \big(a_n c_m + b_n d_m + 2a_n d_m + 2b_n c_m\big)$;

(b) *on each open cell K_k°, s'' is affine, with one-sided end values MX_k^- at x_k and MX_k^+ at x_{k+1}, where*

$$X_k^- = \sum_{n+m=k-1} \Delta a_n \Delta c_m + \sum_{n+m=k} \big(J^F_n \Delta c_m + \Delta a_n J^G_m\big),$$
$$X_k^+ = \sum_{n+m=k} \big(\Delta a_n \Delta c_m + J^F_n \Delta c_m + \Delta a_n J^G_m\big);$$

(c) *on K_k, with s_k the chord of s, $s_k - \frac{1}{8M} \max(0, X_k^-, X_k^+) \le s \le s_k + \frac{1}{8M} \max(0, -X_k^-, -X_k^+)$.*

The correlation $x \mapsto \int F(u)\, G(u - x)\, du$ satisfies the same statements with the knot values $\frac{1}{6M} \sum_{n-m=k} (2a_n c_m + 2b_n d_m + a_n d_m + b_n c_m)$ and second derivative $-\int F'(u)\, G'(u - x)\, du$, the jump of F at 1 included. Consequently, if $0 \le F^- \le f \le F^+$ and $0 \le G^- \le g \le G^+$ on open cells, then the lower bound of (c) for $F^- * G^-$ and the upper bound of (c) for $F^+ * G^+$ enclose $f * g$, and likewise for correlations.

Lemma 5.3 (Signs). *If $F^- \le f \le F^+$ on an open cell, then $f^+ = \max(f, 0)$ lies above the function equal to F^- if both end values of F^- are nonnegative and to 0 otherwise, and below the chord of $\max(F^+, 0)$; and $f^- = \max(-f, 0)$ satisfies the same with $-F^+$ and $-F^-$ in place of F^- and F^+. Products and convolutions of signed functions are enclosed by splitting both operands in this way and combining the four nonnegative terms with their signs.*

Lemma 5.4 (Products and quotients on a cell). *Let p and q be enclosed on a cell by affine bounds $p^\pm$ and $q^\pm$, written in the position $\tau \in [0, 1]$ in the cell, with increments $\Delta p^\pm$, $\Delta q^\pm$ over the cell.*

(a) *For each of the four corner products $r = p^\pm q^\pm$ put $\sigma_r = |\Delta p^\pm \Delta q^\pm|/4$. Then pq lies above the chord through $\min_r(r(0) - \sigma_r)$ and $\min_r(r(1) - \sigma_r)$, and below the chord through $\max_r(r(0) + \sigma_r)$ and $\max_r(r(1) + \sigma_r)$.*

(b) *If either $q^- > 0$ throughout the closed cell or $q^+ < 0$ throughout it, the same holds for $-p/q$, with the corners $r = -p^\pm/q^\pm$ and the slack $\sigma_r = \frac{1}{4}|\Delta q|\,|\Delta p\, q_0 - p_0\, \Delta q|/\min(|q_0|, |q_1|)^3$, where p_0, q_0 and p_1, q_1 are the end values of the corner's bounds; the slack is computed in exact rationals and rounded outward.*

Lemma 5.5 (The forcing chain and the adjoint functionals). *Let the rows be $T_j = A_j \cdot c + B_j h_j + \sum_{i \ge 1} C_{j,i}[h_{j+i}]$ for $1 \le j \le m^*$, where the couplings $C_{j,i}$ are integral operators, and suppose the enclosure of each B_j excludes zero on every closed cell of $[j\eta, \alpha']$ (this is checked).*

(a) *Solving $T_j = 0$ for $j = m^*, \dots, 1$ with the operations of Lemmas 5.1–5.4 and 5.6 encloses the exact forced h_j on every cell.*

(b) *Let a functional have h-part $\sum_m \langle \kappa_m, h_m \rangle$, with $\langle u, v \rangle = \int uv$ over the forcing range (T_0 and $f(1)$ are of this form), and let $\lambda_m = (\kappa_m - \sum_{j<m} C^*_{j,m-j}[\lambda_j])/B_m$ for $m = 1, \dots, m^*$, where $C^*_{j,i}$ is the transpose of $C_{j,i}$. Then $\sum_m \langle \kappa_m, h_m \rangle = -\sum_m \langle \lambda_m, A_m \cdot c \rangle$. Solving the recursion with the same operations encloses the exact λ_m, so both sides are enclosed, the right side without the widths of the forcing chain.*

No approximate solution, and so no residual, enters: each side is an enclosure of an exact quantity, and so is their intersection.

Lemma 5.6 (Blocks). *Let $[\tau_0, \tau_1]$ be a block of thresholds.*

(a) *If an integrand depends on τ only through an indicator $\mathbf{1}_{[1/2-\tau, 1/2)}$ times a nonnegative factor, its integral for τ in the block lies between its values at τ_0 and at τ_1.*

(b) *If $X_\tau(x) = X_{\tau_0}(x) + \int_{\tau_0}^{\tau} r(1/2 - t)\, g(x - 1/2 + t)\, dt$ with $r \ge 0$, then for τ in the block $X_\tau(x) - X_{\tau_0}(x)$ lies between $\rho(\tau) \min g$ and $\rho(\tau) \max g$, the extremes taken over any set containing the arguments $x - 1/2 + t$, $\tau_0 \le t \le \tau$ (the engine takes the cells meeting $[x + \tau - 1/2 - \Delta,\ x + \tau - 1/2]$, with Δ the largest block width), where $\rho(\tau) = \int_{\tau_0}^{\tau} r(1/2 - t)\, dt$. Integrated against a signed H, this bound is applied to H^+ and H^- separately.*

(c) *A quantity depending on σ_β is bounded, for σ_β in a block, by its range over the block.*

The bounds above do not require a rate of continuity. Their additional block losses are $O(\tau_1 - \tau_0)$ in (a) when the integrated density is uniformly bounded, and in (c) when the bounded quantity is uniformly Lipschitz on the block. In (b), if r is uniformly bounded and g is Lipschitz on the argument window, the window-range contribution is $O((\tau_1 - \tau_0)(\tau_1 - \tau_0 + 1/M))$, in addition to the widths of the input enclosures. A window crossing a jump is bounded by its full range; the Lipschitz rate is not asserted there.

Lemma 5.7 (Positivity on closed cells). (a) *If $\Phi \ge L$ on a cell K_n of the large component a, with L affine with end values L_0, L_1, then on K_n*

$$a\, \Phi(a) \ge \min(x_n L_0,\ x_{n+1} L_1) - \max(L_1 - L_0, 0)/(4M).$$

(b) *Suppose $f = f_c + t f_g$, $t \in [t_-, t_+]$, and on the cell $L_c \le f_c$ and $L_g \le f_g \le U_g$, where all three bounds are affine in a. Then $f \ge \min\{L_c + tG : t \in \{t_-, t_+\},\ G \in \{L_g, U_g\}\}$. This minimum is concave in a and lies above the chord through its end values.*

(c) *An open-cell enclosure also bounds the one-sided limits of the enclosed function at that cell's ends. It bounds the assigned value at a grid point if that value is the corresponding one-sided limit. When either adjacent limit can be selected by a class convention, both must be covered. Values on collapsed feasible strata must be checked separately; an almost-everywhere enclosure alone does not bound them.*

For the cases in Table 2, the omitted collapsed free strata have components at extremes of the small interval or at closed upper bin edges. Their constants, parity terms and forced terms are zero, as is checked directly from the class lists and profiles. In the cases with $m^* \geq 2$ the collapsed extended strata with $j_\beta \eta = 1/2$ (all components of β equal to η, and $a = 1/2$; $j_\beta = 9, 10, 15, 15$) also have zero value: every free vector met by the fragmentation of a has at least j_β small components, more than any class with a nonzero constant $(6, 7, 12, 12)$ and more than m^*. At $(1/14, 1/8)$ the stratum $\beta = (1/14, \ldots, 1/14)$, $a = 1/2$ does not have zero value (the constants on $(8, 3, 1)$, $(10, 2, 1)$, $(11, 1, 1)$, $(12, 0, 0)$, $(12, 1, 1)$, $(13, 0, 0)$ enter it); it has the subsum $1/2$ and is not part of (C3). At $(1/11, 1/10)$ there is no such stratum. The other relevant boundary values use the prescribed adjacent class or bin and the continuous limits of the convolution kernels. The excluded half-subsum vectors in the two $m^* = 1$ cases are not part of (C3).

Lemma 5.8 (Binary64 steps)**.** (a) *Let an IEEE 754 binary64 addition, subtraction, multiplication, division or integer conversion have a finite, defined exact result and be rounded to nearest, without overflow. If its returned value is y, its exact result lies between the binary64 neighbours of y.*

(b) *Put $u = 2^{-53}$ and $\gamma_j = ju/(1-ju)$, for $ju < 1$. If k finite binary64 inputs are summed using $k-1$ additions in any order, with no overflow and $(k-1)u < 1$, the absolute error is at most $\gamma_{k-1} \sum_i |x_i|$. More generally, suppose each exact term z_i is first formed from exact binary64 inputs by a product or quotient with at most r rounded operations along its evaluation path. If each of those operations has relative error at most u, and the subsequent sum has no overflow, then, provided $(k-1+r)u < 1$, the absolute error is at most $\gamma_{k-1+r} \sum_i |z_i|$.*

(c) *Let x be a finite binary64 number and $S \geq 0$ an integer with $|x|\, 2^S < 2^{52}$. Then $x\, 2^S$ is exact, and $\lfloor x2^S \rfloor$, $\lceil x2^S \rceil$ are exactly representable integers with $\lfloor x2^S \rfloor \leq x2^S \leq \lceil x2^S \rceil$.*

The separate implementation moves elementwise interval endpoints outward using (a). Its window accumulators also use ordinary rounded divisions, products and additions, followed by aggregate widening using the product-and-sum form of (b), with a reserve for computing the error bound and the final widened endpoints. It convolves the integer images of (c) exactly. The audits checked, on the runs of record, that the scalings are upward, the required values are finite, no operation overflows, the relevant ju are less than 1, every product and quotient in these accumulators is normal, or is zero with a zero operand, or underflows to zero and has as an operand a small integer multiple of 2^{-1074} that is the outward image of an endpoint equal to 0 (a lower endpoint of a nonnegative kernel in a lower accumulator, or an upper endpoint in an upper one, where the kernel vanishes), and the implemented widening factors dominate the required bounds as exact rational inequalities for every length that occurs. A term of the last kind is replaced by 0. In a lower accumulator the true term is nonnegative, so this keeps a lower bound; in an upper accumulator the true term is 0, so the replacement is exact. (A quotient that underflows in this way removes its whole row of terms.) Part (b) is applied to the remaining terms, which are nonnegative. For `p5rig.py`, the preceding nonnegative-accumulator classification applies to the raw updates of `Win.G`. Their operands are the cell enclosures of R_1 and R_2. Lower R_2 endpoints are clipped to zero. A subnormal upper R_2 endpoint in these updates occurs only on a cell where R_2 vanishes; a quotient that rounds it to zero therefore removes a true-zero row. This is a source-specific assertion, not a rule for arbitrary nonnegative inputs. Other interval operations may return nonzero subnormal outward endpoints and are justified by part (a). The signed suffix sums in `kappa_phi` use an absolute-sum error bound: suffixes whose total absolute input sum is subnormal are exact, and otherwise the reserve

in the computed error factor also covers the absolute rounding error of a subnormal error-bound product. For `p4rig3.py`, this classification refers to the raw updates of `Win` at the resolutions of record, and to the initial products and suffix sums in `_xi_kappa`. The used window rows have normal positive endpoints of $\rho = R_1 + R_2$, and their divisions by M are normal. The other raw window operands are the stated nonnegative combinations $Q^{\pm}, U^{\pm}$ of the R_k. Outward additions and scalar multiplications can give negative subnormal lower endpoints; their products round to zero at these resolutions, which is safe because the true kernels are nonnegative. Subnormal upper endpoints in those updates occur where the true combined kernel vanishes, so discarding their underflowed products is also safe. These statements concern the actual kernel sources and resolutions, not arbitrary inputs. Individually widened interval operations may retain nonzero subnormal endpoints and are justified by part (a). Signed functional sums use their absolute-sum error bound and explicit underflow padding. At smaller meshes, signed zero-image artifacts may survive the raw products; the lower artifacts are nonpositive and the upper artifacts are nonnegative on true-zero terms, so retaining them is conservative by monotonicity, while part (b) applies to the remaining normal terms.

5.5. **Independent audits.** By an *independent audit* we mean a review, separate from the implementation, that reads the code bodies against the mathematical text of the case, and recomputes parts of the computation with separately written scripts. The audits did the following.

- $(1/20, 1/5)$, second-order engine: fresh runs at the full resolution $M = 7200$ and at $M = 3600$ reproduce every stored output field exactly. The separate implementation: replays at $M = 36{,}000$ reproduce every printed field; the audit did not rerun the full run. A later dedicated audit of its moving-window code at $1/2$ found that one window mask could omit the assigned value at the right end of a closed window. The consumers of that mask integrate over the endpoint, so the stored bounds remained valid; the audit established the defect and the corrected mask by exact rational checks, replayed $M = 36{,}000$, and bounded the effect on the full run without rerunning it. The author corrected that mask, and a second one in which a further check found a gap that no caller reaches, and reran $M = 36{,}000$ and the full run with the corrected program: every printed line agrees with the earlier outputs apart from timings, and the margin is unchanged. The rerun at $M = 432{,}000$ is the run of record.
- $(1/18, 1/6)$: the exact first-order certificate was replayed at $M = 108{,}000$ and matched every stored field except the elapsed time; the separate implementation's replay at $M = 144{,}000$ reproduces every output byte except timings, and all its integer products at $M = 288{,}000$ were rerun; a later dedicated audit of its window code (see the end of this item) replayed both resolutions end to end with the program as it stood before the corrections described there, reproducing every printed line apart from timings; for the second-order run ($M = 1800$, $NB = 10$) the audit reran the state of the engine it reviewed at $M = 900$, 1800 and 3600, reproducing every field of that state's stored outputs, and the run of record, made by the engine as shipped, equals that $M = 1800$ output in every value field (§10). For the separate implementation, the earlier audit and reproduction did not include a complete code-body sweep of its moving windows at $1/2$. A later dedicated sweep checked the window operands, endpoint conventions, arithmetic and all their consumers. It found that the delta helper's contract, closed in the large component and almost everywhere in the integration parameter, is valid; it must not be read as a joint closed-rectangle bound. It also found a range statement for the stopped U streams that is false outside their used domain, two arithmetic preconditions missing from the program's copy of `vconv.py`, and qualifications needed in the underflow description (now stated after Lemma 5.8). None of these affects the domains that the shipped runs use. The author corrected the program accordingly (the range statement, new assertions, the two guards in `vconv.py` and the header comments;

no numerical path changed) and reran both resolutions with it: every printed line agrees with the earlier outputs apart from timings, with margins ≥ 0.02921 at $M = 144{,}000$ and ≥ 0.04587 at $M = 288{,}000$. The rerun at $M = 288{,}000$ is the run of record.

- $(1/30, 1/4)$: a replay at $M = 1920$ with the run's settings, compared field by field with the stored $M = 7680$ output. The fields recording the settings, inputs, constants and code hashes are identical. Each of the five interval enclosures (of $T_0(c)$, $T_0(g)$, t^*, $f(1)$ and the untwisted T_0 row) contains its $M = 7680$ counterpart. The two positivity bounds and the margin are better at $M = 7680$ (at $M = 1920$ the margin is negative, as expected). This checks consistency across resolutions; it does not reproduce the stored enclosures, and the full run was audited, not rerun. Separately, validated ball-arithmetic integrals that do not use the engine put $T_0(g)$ inside the certified interval and give $0.8705 \leq |B_j| \leq 1.8484$, with sign $(-1)^j$, on the whole forcing range.
- $(1/30, 4/15)$: replays of the engine's components at $M = 960$ (the forcing chain, the parity solve, $f(1)$, the free bound and the adaptively refined extended class), a replay of the diagonals and the parity data at $M = 23{,}040$, and independent validated integrals; the inputs of the full run were checked against the manifest frozen before launch. The full computation was not rerun at any resolution; the extended class that binds the margin was checked by reading its code path.
- Theorem B: its exact-integer cell-convolution computations at $\nu = 3/20$, $1/7$ and $1/8$ (mesh $1/16{,}000$, $1/9800$ and $1/8000$) were made in the audit itself, with separately written scripts in exact integer arithmetic; no exact computation of these three bounds by a second, separately written program is reported here (at $\nu = 3/20$ the same program was also run with mesh $1/8000$, giving a weaker bound). The one at $\nu = 0.116$ (mesh $1/32{,}000$) was made after the audit by a copy of the same program with the term in I_8 added; with that term set to zero the copy reproduces the stored exact outputs at $3/20$, $1/7$ and $1/8$ byte for byte. At $\nu = 0.116$ a second, separately written program, also written and run after the original audit, applies the same scheme with mesh $1/1{,}024{,}000$ and gives a lower bound of at least 0.0354. The scripts and their outputs are in the supplement's `thmB/`.

For $(1/11, 1/10)$ and $(1/14, 1/8)$ there was no independent audit in this sense: the code bodies of the certificates were not audited. A separate review checked their formulas, inputs, signs, normalisations and selected parts of the exact arithmetic against the text, and recomputed in exact arithmetic the least-margin cell of $(1/11, 1/10)$, using that certificate's enclosures of h and $f(1)$. No cell of $(1/14, 1/8)$ was recomputed, and neither certificate was rerun.

The audits' findings (specification fixes, code fixes and wording fixes) were applied before the runs were cited. The audits' separately written recomputations are in the supplement's `audit/` directory, and its README describes, for each case, what the audit or review did, the hash of the certificate's code as run, and its conclusion. The scripts and outputs of the audit of Theorem B, with the computation behind Remark 6.1, are in its `thmB/` directory.

5.6. **Certification status of each case.** We call two implementations *separate* if neither was written from, or with sight of, the other's code. Separate implementations here still share the mathematical text of the case, its input function and the FLINT library for exact integer products. The second-order engine was developed from the code of the first-order exact-integer certificates (their kernel and row formulas, re-implemented with second-order enclosures), so it is not separate from them. In Table 1, "validated binary64" is the separate implementation of §5.2, and the last column counts separate implementations among the rigorous certificates.

For the two cases marked "a float computation agrees", a second computation in binary64 with outward rounding passes. Its rounding control for FFT convolutions rests on an a priori round-off bound for a library routine. We do not treat that as rigorous, so these cases rest on one certificate. For $(1/30, 1/4)$, a second, independent first-order certificate was attempted. Measured on its own

Case (η, α')	Rigorous certificates	Separate implementations
$(1/11, 1/10)$	exact integer, first order; a float computation agrees	1
$(1/14, 1/8)$	exact integer, first order; a float computation agrees	1
$(1/18, 1/6)$	exact integer, first order, and second-order engine (one implementation); validated binary64, first order	2
$(1/20, 1/5)$	second-order engine; validated binary64, first order	2
$(1/30, 1/4)$	second-order engine	1
$(1/30, 4/15)$	second-order engine	1

TABLE 1. Certification status. The base case $(1/11, 1/11)$ of Remark 6.1 is contained in the case $(1/11, 1/10)$ and has no row.

code, it was projected to need M in the millions to reach a positive margin, so it was not run. The cost projection motivates a separately written higher-order implementation; it does not rule out other approaches to a second certificate. No second certificate exists for $(1/30, 1/4)$ or $(1/30, 4/15)$.

6. The cases

Table 2 lists the cases. For each one, the forcing depth m^* is the largest j with $j\eta < \alpha'$ (§3). The constants are the nonzero class constants of the frozen function, used at their exact binary64 values, with $c = 0$ on every other class (Appendix A). The margin is the certified value of $1 - L/U$ (§5), rounded down, at the stated resolution. Theorem 4.1 is applied with m at or below it. The decimal values in this section are outward roundings of the exact enclosures stored with each certificate; each margin fraction is computed from the exact values, not from the printed decimals, and Appendix A gives the exact U and L; the diagonal bounds of the second-order cases are one unit below the logs' nearest-rounded four-decimal values or, for $(1/20, 1/5)$, rounded down from the audit's replay (§10).

6.1. **One small prime per modulus: $(1/11, 1/10)$ and $(1/14, 1/8)$.** Since $\alpha' < 2\eta$, an η-rough modulus $d \le x^\alpha$ is 1 or a prime, and the rows are T_0 and $T_1(\sigma)$ for $\sigma \in [\eta, \alpha']$. The free data are $c[j, k, b] + [j = 1]\, h(s)$ with $h = h_1$. Each row $T_1(\sigma)$ is $A(\sigma) \cdot c + B(\sigma)h(\sigma)$, so $h := -A \cdot c/B$ makes every row with $d = p$ hold exactly. The bins of the largest rough component at $(1/11, 1/10)$ are $(\alpha', 1/4]$, $(1/4, 1/3]$ and $(1/3, 1/2)$; at $(1/14, 1/8)$ there is the single bin $(\alpha', 1/2)$, index 1. The certificates of record give:

- at $(1/11, 1/10)$, $M = 79{,}200$: $B \ge 1.347612$, $T_0(g) \in [-3.600109, -3.599155]$, $t^* \in [-0.002270, 0.004304]$, $f(1) \le -1.095124$, free values at least -1.002270 and extended values at least -1.006102, so the margin is $89022913318324471/1095124185823954082 > 0.08129$;
- at $(1/14, 1/8)$, $M = 107{,}520$: $B \ge 2.091294$, $T_0(g) \in [-3.216895, -3.216306]$, $t^* \in [-0.015790, 0.009699]$, $f(1) \le -1.077374$, free values at least -1.004527 and extended values at least -1.015275, so the margin is $2822701203567884/48971554457620103 > 0.05763$.

The certificates of record are exact-integer, first order, at the resolutions of Table 2. The same code also passes at lower resolutions, with the margins increasing in M: at $(1/11, 1/10)$, with margins at least 0.0278, 0.0611 and 0.0772 at $M = 13{,}200$, 26,400 and 52,800; at $(1/14, 1/8)$, with margin at least 0.0167 at $M = 53{,}760$. The function of record at $(1/14, 1/8)$ has 19 nonzero constants, seven of them on classes of total weight about 10^{-8} that the linear program did not see; the variant with $c = 0$ on those seven classes is certified separately, with margins at least 0.0167, 0.0442 and 0.0576 at $M = 53{,}760$, 80,640 and 107,520. A second computation, in binary64 with outward rounding and FFT convolutions, passes both cases, with the caveat of §5.6.

Label	Case (η, α')	m^*	profile	g	$\lvert c\rvert$	margin	method, M
P3a	$(1/11, 1/10)$	1	all k	$e_{(0,3,3)}$	11	0.08129	exact, 79,200
P3b	$(1/14, 1/8)$	1	all k	$e_{(0,3,1)}$	19	0.05763	exact, 107,520
P4	$(1/18, 1/6)$	2	all k	$e_{(0,3,1)}$	13	0.02909	exact, 216,000
						0.04587	separate, 288,000
P5	$(1/20, 1/5)$	3	$k = 3$	$e_{(0,3,4)}$	32	0.03864	second order, 7200
						0.01609	separate, 432,000*
P6	$(1/30, 1/4)$	7	$w_2 = 1$, $w_3 = -1$	$e_{(0,3,2)} + e_{(0,3,3)}$	45	0.013338	second order, 7680
P7	$(1/30, 4/15)$	7	$w_2 = 1$, $w_3 = -1$	$e_{(0,3,2)} + e_{(0,3,3)}$	49	0.005347	second order, 23,040

TABLE 2. The cases: label (P3 of the computations' records is split into P3a and P3b), forcing depth m^*, profile, parity direction g, number of nonzero class constants $|c|$, and the certified margin with the method and resolution. Classes are written (j, k, b) throughout, with b the bin of the largest rough component ($b = 0$ if $k = 0$); P3b and P4 have the single bin $(\alpha', 1/2)$, index 1. "Exact" is the first-order certificate in exact integer arithmetic and "second order" the second-order engine, in integer arithmetic with directed rounding (§5); "separate" is a separately written first-order certificate in validated binary64 arithmetic, with exact integer convolutions and validated rounding elsewhere (§5.2). Where two margins are given, they come from separate implementations, and both are certificates of record. *Some extended classes on the coarser grid $\{n/72{,}000\}$ (§5.2).

Remark 6.1 (The base case $(1/11, 1/11)$). The region $\nu < 1/11$, $\alpha < 1/11$ lies inside the case $(1/11, 1/10)$ of Theorem A. It also has a simpler construction, which we describe for illustration only: Theorem 4.1 assumes $\eta < \alpha'$, and nothing in this paper rests on the construction. With $\eta = \alpha' = 1/11$ no forced function occurs, and every modulus $1 < d \le x^\alpha$ with $\alpha < 1/11$ has a prime factor at most $x^\alpha < (x/2)^\eta$ for large x (this uses $\alpha < 1/11$ strictly), so it divides no element of the support of the perturbation; the only twisted row is T_0. The free data depend only on the number k of components, $f = c_k$: the free values $(-1)^k$ of Ford and Maynard's $\tilde{\lambda}^{(1/2,\eta)}$ [5, §9.1] (compare the multiplicative ansatz of [3, (3.8)]) replaced by general coefficients. Lemma 2.1 reduces the extension to the one-variable functions $J_k(a) = a \int_{\eta \le u_1 < \cdots < u_k,\ u_1 + \cdots + u_k = a} L_{1/2}(u)\, \frac{du}{u_1 \cdots u_k}$. The constants c_3, c_4 are solved from $f(1) = -1$ and $T_0 = 0$, with $c_5, \ldots, c_{11}$ fixed; since $J_k(1) = 0$ for even k ($L_{1/2}$ vanishes almost everywhere on vectors of even length and sum 1, by Lemma 2.1), c_4 does not enter $f(1)$. At $\eta = 1/11$ an exact-integer computation made during the independent audit ($M = 4400$) gives $3.31449 < T_0(e_4) < 3.38132$, where e_4 is the unit vector of c_4, encloses (c_3, c_4) in a box, and shows over that box that $f \ge -0.87595$ on free and $f \ge -0.89904$ on extended non-prime vectors $u \ne E$, with $f(1) = -1$; at $\eta = 1/12$ the bounds are -0.75540 and -0.77919.

6.2. **Two small primes: $(1/18, 1/6)$.** Here $m^* = 2$, and $3\eta = \alpha'$: level 3 is a null level, $h_3 \equiv 0$, and moduli with three small primes exceed x^α only because $\alpha < 1/6$ strictly. The free data are $c[j, k, b] + [j = 1]\, h_1(s) + [j = 2,\ s_1 + s_2 \le 1/6]\, h_2(s_1 + s_2)$, with the single bin. The rows at $d = p_1 p_2$ force $h_2 = -A_2 \cdot c/B_2$, with $B_2 < 0$; the rows at $d = p$ then force h_1. The thirteen constants come from a linear program. The case has three audited certificates, from two separate implementations (§5.6): exact first order (≥ 0.0067, 0.0179, 0.0290 at $M = 108{,}000$, 144,000, 216,000); the second-order engine (≥ 0.0546 at $M = 1800$, $NB = 10$), developed from the first-order code and so not separate from it; and the separate implementation in binary64 interval arithmetic with exact integer convolutions (≥ 0.02921 at $M = 144{,}000$ and ≥ 0.04587 at $M = 288{,}000$). The three enclose different conservative quantities: the first bounds the couplings of h_2 over fixed blocks (§5). At $M = 216{,}000$ the exact first-order certificate gives $B_1 \ge 1.306391$, $B_2 \le -1.306391$, $T_0(g) \in [-2.524087, -2.523870]$, $t^* \in [-0.012289, 0.017693]$, $f(1) \le -1.049814$, free

values at least -1.003538 and extended values at least -1.019267; its margin is $30548429835746437/1049814788771959394 > 0.02909$, and the lift uses $m = 0.029$.

6.3. **Three small primes, one rough count:** $(1/20, 1/5)$**.** Here $m^* = 3$, and $4\eta = \alpha'$: four small components have total at least α', so h_4 could live only at the single total α', a null set, and $h_4 \equiv 0$. Moduli with four small primes exceed x^α because $\alpha < 1/5$ strictly. The forced functions are carried only on vectors with exactly three rough components (the profile $K = \{3\}$), and the bins are $(1/5, 1/4]$, $(1/4, 1/3]$, $(1/3, 5/12]$, $(5/12, 1/2)$. The certified diagonals satisfy $|B_j| \geq 0.825079$ with alternating signs, $T_0(g) \in [-1.077245, -1.077243]$, $t^* \in [-0.012088, -0.007984]$, $f(1) \leq -1.055341$, the free minimum is at least -1.014557 and the extended minimum at least -1.011449. The margin is at least $40784885613599943/1055341401115063442 > 0.03864$ at $M = 7200$, and the lift uses $m = 0.038$. The separate implementation certifies the case at $M = 432{,}000$ with margin at least 0.01609; some of its extended classes are bounded on the coarser grid $\{n/72{,}000\}$ (§5.2).

6.4. **Seven small primes, weighted profile:** $(1/30, 1/4)$ **and** $(1/30, 4/15)$**.** Here $m^* = 7$. The profile is weighted, $w_2 = 1$ and $w_3 = -1$, so the diagonal is the weighted one of Lemma 3.1. The weights matter: the two-rough part alone has a minimum of about 0.14 at $(1/30, 1/4)$, while the weighted diagonal is at least 0.8704. The bins of the largest rough component are $(\alpha', 1/3]$, $(1/3, 5/12]$, $(5/12, 1/2)$ in both cases.

- At $(1/30, 1/4)$, $8\eta > \alpha'$, so there is no level 8. The certificate at $M = 7680$ gives $|B_j| \geq 0.8704$ at every level, $T_0(g) \in [-0.883325, -0.883323]$, $t^* \in [-0.009130, 0.015907]$, $f(1) \leq -1.02811366$, free values at least -1.01440047 and extended values at least -1.01192797. The margin is
$$2742639303632839/205622732434150043 > 0.013338,$$
and the lift uses $m = 0.013$.
- At $(1/30, 4/15)$, $8\eta = \alpha'$: level 8 is a null level, $h_8 \equiv 0$, and moduli with eight small primes exceed x^α only because $\alpha < 4/15$ strictly. The certificate at $M = 23{,}040$ gives $|B_j| \geq 0.5669$ at every level, with signs $(-1)^j$, $T_0(g) \in [-0.556040, -0.556039]$, $t^* \in [0.013941, 0.033589]$, $f(1) \leq -1.02099183$, free values at least -1.01477745 and extended values at least -1.01553199; the extended class $(0, 1)$ binds. The margin is
$$363989972040679/68066122248838481 > 0.005347,$$
and the lift uses $m = 0.005$. On $\sigma \geq 1/5$ the three-rough part of the diagonal vanishes, and the remaining two-rough part satisfies $|B_j| \geq 1.277$. On the whole forcing range the certificate gives the uniform lower bound $|B_j| \geq 0.5669$; this is an enclosure, not a claim that the value is attained. The lift uses this separation from zero to choose ε small enough that the masked diagonals remain nonzero.

In both cases the free bound is the superset bound of §5, labelled $(5, 2, 1)$. At $(1/30, 1/4)$ no vector of that class carries h (its small totals lie in $[1/3, 1/2)$), so the bound is conservative; the floating-point candidate for the actual minimum is about -1.0120 (not certified). At $(1/30, 4/15)$ the bound equals that of the realisable class $(5, 2, 3)$.

7. Parity forces primes from $\nu = 0.116$

This section proves Theorem B. We work with Ford and Maynard's parameters $P = (1/2, 0, \nu)$, except in Lemma 7.1 and in the derivation of $(*)$, which allow general parameters $P = (\gamma, \theta, \nu)$; the objects below are taken for the parameters in use. We use their notation: $|z|$ and $\dim z$ for the sum and the number of components of a vector z, and z_A for its subvector indexed by $A \subseteq [\dim z]$ [5, Definitions 4.1 and 4.2]; $[k] = \{1, \dots, k\}$ and $\Omega(n)$ [5, §5.1]; $R = R(P)$ for the fundamental region and $C(R)$ for the set of coagulations of vectors of R [5, Definitions 4.3 and 4.5]; $v(n)$ and $v(m; n)$ for the vectors of the numbers $\log p / \log n$ over the prime factors p of n, respectively of m, counted

with multiplicity and in increasing order [5, Definition 5.1]; $P^-(n)$ and $P^+(n)$ for the smallest and the largest prime factor of n, with $P^-(1)=\infty$ and $P^+(1)=1$ [5, §5.1]; $\mathcal{P}$ for the set of primes in $(x/2,x]$ and $n\sim x$ for $x/2<n\le x$ [5, §7]; $\mathcal{H}(P)=\{z\in C(R(P)):\dim z\ge 2,\ z_i\ge\nu \text{ for all } i\}$; G_1 for the set of admissible weight vectors [5, (7.1)]; "admissible g" for a function in their class of weight functions supported on G_1 [5, Definition 7.2]; $\mathcal{N}=\{x/2<n\le x : n \text{ composite},\ v(n)\in C(R)\}$ [5, (7.12)]; and, for g on G_1, the ordinary-divisor weight $A_g(n)$ of [5, (7.13)–(7.14)] and the vector convolution $V_g=1\star g$ [5, Definition 7.1]. Numbered items cited from [5] carry Ford and Maynard's numbering; Lemma 7.1, Proposition 7.2 and Lemma 7.3 are this paper's.

On squarefree $n\in\mathcal{N}$ with $P^-(n)\ge n^\nu$, the ordinary-divisor weight $A_g(n)$ agrees with $V_g(v(n))$. On non-squarefree n the weights can differ, because ordinary divisors are counted once while the vector convolution counts subvectors by index. For example, take $P=(1/2,0,\nu)$ with $0<\nu<1/4$ and an admissible g with $g(\emptyset)=1$ and $g((1/2))=-1/2$; its support forces it to vanish on vectors of sum 1. If $x/2<p^2\le x$, then $p^2\in\mathcal{N}$ and $P^-(p^2)>(p^2)^\nu$, but $A_g(p^2)=1+g((1/2))=1/2$, whereas $V_g((1/2,1/2))=1+2g((1/2))=0$. Ford and Maynard identify the ordinary-divisor weight with the corresponding vector convolution in the second sentence after [5, (7.14)], beginning "The above definition of H is equivalent", and, with a roughness indicator, in [5, Lemma 7.18(a)] and the final sentence of the proof of part (a). The same counting issue occurs in the expansion over $J\subseteq[k]$ at the start of the proof of [5, Proposition 7.22], there for $n\notin\mathcal{P}\cup\mathcal{N}$. For our application we need [5, (7.15)] for the ordinary-divisor weight, which is how it is stated; its printed proof needs corrections, one of them at repeated prime factors, and Lemma 7.1 gives the argument in full. It follows the steps of the proofs of [5, Lemma 7.21 and Proposition 7.22], the latter with their parameters $\sigma=\nu$ and $\lambda(d)=\mu(d)$.

Lemma 7.1. *Let (γ,θ,ν) satisfy Ford and Maynard's standing conditions [5, (6.1)] with $\gamma\ge 1/2$ and $\nu\le 1-\gamma$, let g be admissible, and let $A_g(n)=\sum_{d\mid n}G(d;n)$ be the weight of [5, (7.13)–(7.14)]. Let $A\ge 1$ and $\varpi\ge 1$. If B is large enough in terms of $A,\varpi,\gamma,\theta,\nu$ and g, and (w_n) satisfies (w), (I) and (II), then $\sum_{n\in\mathcal{P}\cup\mathcal{N}}w_nA_g(n)\ll_{g,A}x(\log x)^{-A}$.*

Proof. All constants may depend on γ,θ,ν and g, and x is taken large in terms of these and of the logarithmic exponents below. For $d\le x^\gamma$ put

$$\Delta_d=\sup_{\text{interval } I}\Big|\sum_{\substack{dm\sim x\\ m\in I}}w_{dm}\Big|,$$

so that (I) gives $\sum_{d\le x^\gamma}\tau(d)^C\Delta_d\le x(\log x)^{-B}$ for every $C\le B$. For a bounded function h on vectors, write $A_h(n)=\sum_{d\mid n}G_h(d;n)$, with G_h defined as G in [5, (7.13)] with h in place of g.

Boundary pieces. Let h be bounded and supported on vectors of a fixed dimension $j\ge 1$ that lie on a proper affine hyperplane $c_1t_1+\cdots+c_jt_j=c_0$, $(c_1,\dots,c_j)\ne 0$. A term of $A_h(n)$ comes from $d=m_1m_2\mid n$, $d\le n^\gamma$, with $m_2=p_1\cdots p_j$, $p_i\ge n^\nu$, and $\sum_i c_i\log p_i=c_0\log n$. If $c_0\ne 0$, then m_2 determines n, so for each $d\le x^\gamma$ at most $\tau(d)$ integers $n\sim x$ receive a term from d, and each $|w_n|$ is at most Δ_d (take I to be the single cofactor). Hence $\sum_{n\sim x}|w_nA_h(n)|\ll\sum_{d\le x^\gamma}\tau(d)\Delta_d$. If $c_0=0$, let $\mathcal{D}$ be the set of $m_2\le x^\gamma$ whose j prime factors are at least $P_0=(x/2)^\nu$ and satisfy $\sum_i c_i\log p_i=0$. Choosing i with $c_i\ne 0$, the other $j-1$ primes determine p_i, so $\sum_{m_2\in\mathcal{D}}1/m_2\ll P_0^{-1}(\log x)^{j-1}$. With $S_d=\#\{m_2\mid d : m_2\in\mathcal{D}\}$ and $\rho(n)=\sum_{d\mid n,\,d\le x^\gamma}S_d$ we have $|A_h(n)|\ll\rho(n)$,

$$\sum_{n\sim x}\rho(n)\le x\sum_{d\le x^\gamma}\frac{S_d}{d}\ll x^{1-\nu}(\log x)^j,\qquad \Big|\sum_{n\sim x}\rho(n)w_n\Big|\le\sum_{d\le x^\gamma}\tau(d)\Delta_d,$$

and, since $w_n\ge -x^{\nu/10}$ by (w), $\sum_{n\sim x}|w_n|\rho(n)\le|\sum_n\rho(n)w_n|+2x^{\nu/10}\sum_n\rho(n)$. In both cases $\sum_{n\sim x}|w_nA_h(n)|\ll x(\log x)^{-B}+x^{1-9\nu/10}(\log x)^j$, and the bound holds for the sum over any set of n.

Reduction to Lipschitz pieces. Write the restriction of g to ordered vectors as the finite sum of functions g_a on polytopes Q_a given by its definition, and work with each summand separately. In ambient dimension $j \geq 1$, intersect Q_a with $\{t_i \geq \nu,\ \sum_i t_i \leq \gamma\}$, and call the intersection P_a. This does not change any contributing term of A_{g_a}. If P_a is full-dimensional, the bounded first derivatives of g_a on its interior give a bounded Lipschitz function there, which extends continuously to its closure. Restrict this extension to P_a. Its difference from $g_a \mathbf{1}_{P_a}$ is supported on the boundary of P_a, which is contained in a finite union of proper affine hyperplanes. If P_a is not full-dimensional, including when it is a point in $\mathbb{R}^j$ with $j \geq 1$, the whole piece is supported on a proper affine hyperplane. The boundary estimate treats all these differences in absolute value, over any set of integers. Keep the value $g(\emptyset)$ separately. By linearity it therefore suffices to handle a bounded Lipschitz function $\tilde{g}$ on a convex polytope $\mathcal{G}$, multiplied by its indicator, and the term at $\emptyset$.

The full sum. By linearity fix one Lipschitz piece and fix m_1, m_2 with $m_1 m_2 \leq x^\gamma$. Write $m_2 = q_1 \cdots q_j$ with prime factors counted with multiplicity; for $m_2 = 1$ this is the empty list. The conditions $m_1 m_2 \leq n^\gamma$, $P^+(m_1) < n^\nu \leq P^-(m_2)$ and $v(m_2; n) \in \mathcal{G}$ restrict $n \sim x$ to an interval, since

$$v(m_2; n) = (\log q_1, \dots, \log q_j)/\log n$$

moves along a ray. On this interval the Lipschitz bound gives $|\tilde{g}(v(m_2; t)) - \tilde{g}(v(m_2; t'))| \ll |t - t'|/(x \log x)$. Its total variation, including the bounded endpoint jumps when extended by zero, is therefore $O(1)$. Partial summation within the interval and (I) give

$$\sum_{x/2 < n \leq x} w_n A_{\tilde{g}\mathbf{1}_{\mathcal{G}}}(n) \ll \sum_{d \leq x^\gamma} \tau(d) \Delta_d \leq x(\log x)^{-B}$$

for each reduced Lipschitz piece $(\tilde{g}, \mathcal{G})$, and likewise for the term at $\emptyset$. This is the Type I argument of [5, Lemma 7.21]. It remains to bound their sum over $n \notin \mathcal{P} \cup \mathcal{N}$.

The complementary sum: factorisation and polytopes. Write uniquely

$$n = a p_1 \cdots p_k, \qquad P^+(a) < n^\nu \leq p_1 \leq \cdots \leq p_k, \qquad 0 \leq k \leq 1/\nu.$$

This is the factorisation of [5, (7.13)]. Let $\mathcal{U}_1, \dots, \mathcal{U}_{J_0}$ be the complement polytopes constructed below, in positive ambient dimension, for small vectors with positive coordinates less than ν, including the empty vector. For $n \sim x$ with $\Omega(n) \geq 2$, the condition $v(n) \notin C(R)$ is equivalent to $k = 0$ or $v(p_1 \cdots p_k; n) \in \bigcup_j \mathcal{U}_j$. For $k = 0$, all coordinates of $v(n)$, and hence of every fragmentation of $v(n)$, are less than ν; no such fragmentation lies in R [5, Lemma 5.6(i)], so $v(n) \notin C(R)$. For $k \geq 1$, use the form verified below with the small vector $v(a; n)$. In particular, if $t \in \mathcal{U}_j$ and q is any vector, possibly empty, of positive coordinates less than ν with $|q| = 1 - |t|$, then $(t, q) \notin C(R)$. We use [5, Lemma 5.8] only for small vectors y whose coordinates are positive and less than $\sigma = \nu$, possibly $y = \emptyset$, and we check this form directly, following its proof. As there, $(y, x) \in C(R)$ if and only if x is a coagulation of a vector z of bounded dimension with positive coordinates and $(y, z) \in R$. If $y \neq \emptyset$, then for each $J \subseteq [\dim z]$ the sums $|z_J| + y_1 + \cdots + y_i$ $(0 \leq i \leq \dim y)$ increase from $|z_J|$ to $|z_J| + |y|$ in steps less than ν; using $y_1 < \nu$ when $J = \emptyset$, and $\theta + \nu < 1$ when $J = [\dim z]$, it follows that $(y, z) \in R$ if and only if $|y| + |z| = 1$, $z \in (0, 1 - \gamma)^{\dim z}$, and $|z_J| > \theta + \nu$ or $|z_J| < \theta - |y|$ for every $J \subseteq [\dim z]$, a condition on $|y|$ alone. If $y = \emptyset$, the same holds with J restricted to $\emptyset \neq J \neq [\dim z]$, since the empty and the full subsum are not proper. With $\mathcal{Q}_\ell(\mathcal{J})$ as in the proof of [5, Lemma 5.8], define $\mathcal{Q}^0_\ell(\mathcal{J})$ in the same way using only the proper nonempty subsets J of $[\ell]$. For $1 \leq k \leq \lfloor 1/\nu \rfloor$ and $k \leq \ell \leq k + \lfloor 2/(1-\gamma) \rfloor$, fix an ordered partition $\pi = (I_1, \dots, I_k)$ of $[\ell]$ into nonempty parts and write $T_\pi(z) = (|z_{I_1}|, \dots, |z_{I_k}|)$. Each of the sets

$$T_\pi(\mathcal{Q}_\ell(\mathcal{J})) \cap [\nu, 1]^k \cap \{|x| < 1\}, \qquad T_\pi(\mathcal{Q}^0_\ell(\mathcal{J})) \cap [\nu, 1]^k \cap \{|x| = 1\}$$

is a convex polytope with boundedly many constraints, each strict or not: T_π is linear, and eliminating the coordinates of z (Fourier–Motzkin elimination, which applies to strict and non-strict inequalities alike) describes the image by boundedly many such constraints, as at the end of the

proof of [5, Lemma 5.8]. Take their union over the finitely many choices of ℓ, $\mathcal{J}$ and π. Its intersection with $\{|x| = 1 - |y|\}$ is exactly the set of x with $x_i \geq \nu$ and $(y, x) \in C(R)$: the first family applies when $y \neq \emptyset$, and the second when $y = \emptyset$. For each k, subdivide the complement of this union in $[\nu, 1]^k \cap \{|x| \leq 1\}$ into disjoint convex polytopes. These are the $\mathcal{U}_j$. The number of pieces and of their defining constraints is bounded in terms of the fixed parameters.

We split the sum according to the polytope, to the set $\mathcal{W} = \{i : p_i = p_{i+1}\}$, and to the term of A_g: if J runs over the subsets of $[k]$ that meet each maximal block of equal primes in an initial segment, then $J \mapsto \prod_{j\in J} p_j$ is a bijection onto the divisors of $p_1 \cdots p_k$, and

$$A_g(n) = \sum_{u|a} \sum_{J} \mathbf{1}\big(u \textstyle\prod_{j\in J} p_j \leq n^\gamma\big)\, \mu(u)\, g\big(v(\prod_{j\in J} p_j; n)\big).$$

There are $O(1)$ choices of k, $\mathcal{W}$, J and a piece $(\tilde{g}, \mathcal{G})$; the term at $\emptyset$ is the piece $\mathcal{G} = \{\emptyset\}$, $\tilde{g} \equiv g(\emptyset)$, which occurs only with $J = \emptyset$, and a piece in $\mathbb{R}^j$ with $j \geq 1$ occurs only with $|J| = j$. For $k \geq 1$ and each such choice, let $\mathcal{T}'$ be the intersection of the selected complement polytope with $\{\nu \leq t_1 \leq \cdots \leq t_k\}$, with $t_i = t_{i+1}$ $(i \in \mathcal{W})$ and $t_i < t_{i+1}$ $(i \notin \mathcal{W})$, and with $\{t : t_J \in \mathcal{G}\}$. It is a convex polytope with boundedly many constraints, each strict or not, and it keeps the property just stated. For each $u \mid a$ write $v = a/u$, so that $a = uv$. For $k \geq 2$ the condition $\Omega(n) \geq 2$ is automatic; for $k \leq 1$ it is $uv > 1$, which we keep. For $k = 0$ there is one empty equality pattern, $J = \emptyset$, and the factor $g(\emptyset)$ is constant; no $\mathcal{T}'$ condition is imposed.

In the reductions below, fixed bounds for the separate coefficient functions are absorbed into the implied constants, so that these functions may be taken 1-bounded. When variables are grouped to apply [5, Lemmas 7.9 and 7.17], the coefficient arrays are bounded by $|w_n|\tau(n)^C$ for a fixed C, and are set to zero off $n \sim x$. By (w) and the divisor bound,

$$|w_n|\tau(n)^C \leq x(\log x)^{\varpi}\tau(n)^C < x^2$$

for sufficiently large x. Thus their coefficient hypotheses hold.

Prime decomposition ($k \geq 1$). We apply the box and prime decomposition argument of [5, Lemma 7.14] to the tuples with $P^+(uv) < n^\nu \leq p_1$, taking the first logarithmic box interval for u and v closed at zero, to include unit factors. The argument needs only the following: for fixed (u, v) the conditions define a convex region in the coordinates $\log p_i / \log x$; $p_i \geq n^\nu > (x/2)^\nu$ gives the lower bound used there; within each box the Lipschitz bound for $\tilde{g}$ applies, since $\mathcal{T}' \subseteq \{t : t_J \in \mathcal{G}\}$, and, as $n > x/2$, each $\log p_i / \log n$ varies over a box by at most a constant times its side, also with the first interval closed at zero; the factorisation of n is unique, so the error of the box approximation is $O((\log x)^{-D} \sum_n |w_n|\tau(n))$ with a fixed $D \geq A + \varpi + 2$, taking B sufficiently large after D is chosen; and [5, Lemma 7.12] applies to these regions, including the faces given by $\mathcal{W}$; with the first box interval closed at zero, the box condition for u or v is a one-sided bound $u \leq x^c$ or $v \leq x^c$, which [5, Lemma 7.9] removes as in the proof of [5, Lemma 7.13]. From here on $\ell = 6\lceil 1/(1-\gamma)\rceil$, as in [5, Lemma 7.12]. In applying the prime decomposition here and in §7.5, we make the following cutoff explicit. In the proof of [5, Lemma 7.12] (in the notation of that proof), retain $f_{h,j}^{r_h - 1} \leq x$ when separating the identity $d_{h,j} = f_{h,j}^{r_h-1}$. This cutoff is automatic before separation, since $d_{h,j} \mid n \leq x$, and is absorbed into the bounded coefficient of $f_{h,j}$ afterwards. On its support the compared quantities $d_{h,j}$ and $f_{h,j}^{r_h-1} \pm \frac{1}{2}$ lie in $[1/2, x + 1/2]$, and their distance is at least $1/2$. Off its support, use the constant functions $1/2$ and $3/2$ in place of $f_{h,j}^{r_h-1} - 1/2$ and $f_{h,j}^{r_h-1} + 1/2$, respectively; the coefficient there is zero. The coefficient arrays there are sums over the remaining variables, so they are at most $\max_{n\leq x} \tau_K(n) \sum_n |w_n|\tau(n) < x^2$ for a fixed K and large x, by (w) and the divisor bound. Thus the hypotheses of [5, Lemma 7.9] hold uniformly in r_h, and the logarithmic losses and errors in the prime decomposition are unchanged. The argument reduces the estimate, at the cost of a fixed power of $\log x$ and an error $O(x(\log x)^{-A-1})$, to sums over $n = uvm_1 \cdots m_k \sim x$ with $m_h = d_{h,1} \cdots d_{h,\ell}$, $d_{h,j} < n/x^\gamma$, bounded coefficients depending on

one $d_{h,j}$ each, the coefficient $\mu(u)u^{it}v^{it'}$ for real t, t', and the conditions

$$\tilde{y} = \big(\log m_h/\log n\big)_{h\le k} \in \mathcal{T}', \qquad P^+(uv) < n^\nu, \qquad u\prod_{h\in J} m_h \le n^\gamma.$$

For $k = 0$ we skip this step: the sum is already a sum over $n = uv$ with the coefficient $\mu(u)$.

Splitting the small part. Take $L = x^{\nu/3}$; for large x, $2 \le L \le (x/2)^{\nu/2}$. We apply the smooth splitting of [5, Lemma 7.17] in the form with $P^+(d) < (dm)^\nu$ in the input and $d_j < (dm)^\nu$ for every block in the output. For this form of the smooth split, a composite canonical block d_j satisfies $d_j \le LP^-(d_j) \le Ld_j^{1/2}$, so $d_j \le L^2 \le (x/2)^\nu < (dm)^\nu$, while for a prime block the two conditions coincide. By [5, Lemma 7.16], $\mu(u)u^{it}$ and $v^{it'}$ are splittable with respect to L. Applying this form to u and then to v gives, up to a fixed power of $\log x$ and $O(x^{-80})$, sums with $u = u_1\cdots u_s$, $v = v_1\cdots v_r$, $r, s \le 3/\nu + 2$, a bounded coefficient on each block, and

$$u_i, v_j < n^\nu, \qquad u_i > x^{\nu/3}\ (i < s), \qquad v_j > x^{\nu/3}\ (j < r), \qquad u_s, v_r \le x^{\nu/3}.$$

Let $\mathcal{I}$ be the indicator of the coarse conditions

$$u_i > x^{\nu/3}\ (i < s), \qquad v_j > x^{\nu/3}\ (j < r), \qquad u_s, v_r \le x^{\nu/3}, \qquad m_h > (x/2)^\nu\ (h \le k).$$

The last conditions follow from $\tilde{y}_h \ge \nu$. Retain $u_i < n^\nu$ for $i < s$ and $v_j < n^\nu$ for $j < r$ as separate summation conditions. The corresponding bounds for u_s, v_r follow from their coarse upper bounds and $n > x/2$, for large x.

A Type II factor. Form z by taking $\log b/\log n$ for every entry b of the indexed list

$$(u_1, \dots, u_s, v_1, \dots, v_r, (d_{h,j})_{h,j}),$$

omitting entries equal to 1. Equal entries at different indices are retained. The coordinates of z are positive, sum to 1, and are less than $1-\gamma$, since $u_i, v_j < n^\nu \le n^{1-\gamma}$ and $d_{h,j} < n/x^\gamma \le n^{1-\gamma}$. Let q be the vector of $\log u_i/\log n$ and $\log v_j/\log n$ for those $u_i, v_j > 1$. Its coordinates are less than ν and $|q| = 1 - |\tilde{y}|$. If $k \ge 1$ and $z \in R$, joining the d-coordinates belonging to each m_h gives $(\tilde{y}, q) \in C(R)$, contrary to the property of $\mathcal{T}'$. Hence, for $k \ge 1$, $z \notin R$, so it has a nonempty proper subsum in $[\theta, \theta+\nu]$. If $k = 0$, every coordinate of z is less than ν; adding coordinates until the sum reaches θ (or taking one coordinate if $\theta = 0$) gives a positive subsum in $[\theta, \theta+\nu]$, proper since $\theta + \nu < 1$. In both cases the subsum lies in

$$\mathcal{K}_n^* = \Big(\theta\frac{\log(x/2)}{\log n},\ (\theta+\nu)\frac{\log x}{\log n}\Big],$$

because $n > x/2$ when $\theta > 0$, the subsum is positive when $\theta = 0$, and $n \le x$. Label all blocks, units included, as $n_1, \dots, n_K$, where $n_i = u_i$ for $1 \le i \le s$, $n_{s+j} = v_j$ for $1 \le j \le r$, followed by the $d_{h,j}$ in lexicographic order; thus $K = k\ell + r + s$. For nonempty $E \subseteq [K]$, $\sum_{e\in E} \log n_e/\log n \in \mathcal{K}_n^*$ is equivalent to $(x/2)^\theta < \prod_{e\in E} n_e \le x^{\theta+\nu}$, the range of (II), and this excludes products of units. By inclusion–exclusion over the nonempty subsets of $[K]$ it suffices to treat sums with these conditions for every E in a fixed nonempty collection $\mathcal{E}$. We remove each condition $d_{h,j} < n/x^\gamma$, written as $n/d_{h,j} > \lfloor x^\gamma\rfloor + 1/2$, by [5, Lemma 7.9]; the twists it introduces depend on one block each and are absorbed into the coefficients.

Removing the polytope conditions. We treat $uv = 1$ and $uv > 1$ separately. If $uv = 1$, then every u_i and v_j is 1, these unit conditions are absorbed into the coefficients, $\mathcal{I}$ forces $r = s = 1$, and $k \ge 2$, since for $k \le 1$ we kept $uv > 1$. On both branches we keep every constraint of $\mathcal{T}'$, the cutoff $u\prod_{h\in J} m_h \le n^\gamma$, and the strict bounds $u_i < n^\nu$ $(i < s)$ and $v_j < n^\nu$ $(j < r)$ as separate conditions to be removed by [5, Lemma 7.11]. Its designated blocks are the sets of variables whose products are the m_h, and the single variables u_i $(i < s)$ and v_j $(j < r)$; on $\mathcal{I}$ each has product at least $x^{\nu/3}$, as that lemma requires with the ν of (w). At least one exists, since $k = 0$ and $r = s = 1$ would give $n \le x^{2\nu/3}$. Each condition is a monomial inequality $n_1^{c_1}\cdots n_K^{c_K} \le 1$ or < 1 with equal nonzero exponents on some designated block, as follows. For $i < s$, the condition $u_i < n^\nu$ is

$n_i^{1-\nu}\prod_{e\neq i} n_e^{-\nu} < 1$, with the designated singleton $\{i\}$; the same argument applies to each v_j with $j < r$ using its own index. The cutoff has exponent $1-\gamma$ on the variables of u and of m_h ($h \in J$) and $-\gamma$ on the others, and every block is wholly of one kind. A linear constraint $\sum_h e_h\tilde{y}_h \le e_0$ (or $< e_0$) of $\mathcal{T}'$ reads $(uv)^{-e_0}\prod_h m_h^{e_h-e_0} \le 1$ (or < 1); if $e_h \neq e_0$ for some h, the block of m_h serves. If $e_h = e_0$ for every h, the constraint is a condition on $|\tilde{y}|$ alone, and $|\tilde{y}| = 1$ when $uv = 1$, $|\tilde{y}| < 1$ when $uv > 1$, so on each branch it either always holds, and is dropped, or never holds, and the summation piece is empty. The constraints from $\mathcal{W}$ have $e_0 = 0$ and nonzero exponents on the blocks of m_i and m_{i+1}; an equality is written as two weak inequalities. After a bounded positive rescaling the exponents lie in $[-D', -1]\cup\{0\}\cup[1, D']$ for a fixed D'. We apply [5, Lemma 7.11] with $\phi_{\mathbf{n}}$ equal to the product of the coefficients, $\mathcal{I}$, the branch condition and the conditions indexed by $\mathcal{E}$; both sums it produces have the same form, and its error term is $O(x^{1-\nu/20})$. The block conditions $\prod_{e\in\mathcal{M}} n_e \ge x^{\nu/3}$, one for each designated block $\mathcal{M}$, that remain in those sums are implied by $\mathcal{I}$, which is part of $\phi_{\mathbf{n}}$, and are dropped.

Type II. Every condition in $\mathcal{I}$ now compares an integer product Π with a fixed threshold Ξ. Write a strict lower condition as $\Pi > \lfloor\Xi\rfloor + 1/2$ and a weak upper condition as $\Pi \le \lfloor\Xi\rfloor + 1/2$. Apply [5, Lemma 7.9] to these conditions and, on the second branch, to $uv > 1$, written as $u_1\cdots u_s v_1\cdots v_r > 3/2$. The unit indicators on the first branch remain in the separate coefficients. If $[K] \in \mathcal{E}$, the whole summation piece is empty for large x, since $n > x/2 > x^{\theta+\nu}$, and is discarded. Otherwise every member of $\mathcal{E}$ is proper. Choose one $E \in \mathcal{E}$ and remove the interval conditions indexed by the other members of $\mathcal{E}$ using the same half-integer thresholds. The gaps are at least $1/2$ and both compared quantities lie in $[1/2, 2x]$, as [5, Lemma 7.9] requires, and all twists factor into one-variable coefficients. The remaining sum is

$$\sum_{\substack{n'n''\sim x\\ (x/2)^\theta < n' \le x^{\theta+\nu}}} Y_1(n')Y_2(n'')w_{n'n''},$$

where Y_1 and Y_2 sum the coefficients over the factorisations of n' into the blocks of E and of n'' into the others, so that $|Y_1(n')| \le \tau(n')^K$ and $|Y_2(n'')| \le \tau(n'')^K$. By (II) it is at most $x(\log x)^{-B}$ once $B \ge K$. All losses were fixed powers of $\log x$ and power savings, so for B large in terms of A and ϖ the complementary sum is $O(x(\log x)^{-A})$. Subtracting the complementary sum from the full sum, and restoring the boundary pieces by their absolute-value estimate, proves the lemma for the original g. □

We use Lemma 7.1 for A_g and, separately for the vector weight V_g, condition (b.2) as in the proof of [5, Lemma 7.20]; Proposition 7.2 imposes an arithmetic sign condition on non-squarefree $n \in \mathcal{N}$ with $P^-(n) \ge n^\nu$ outside an exceptional set with negligible absolute w-mass.

For any parameters $P = (\gamma, \theta, \nu)$, let g be admissible and $n \in \mathcal{N}$. As in the proof of [5, Lemma 7.18(a)], the cutoff $d \le n^\gamma$ in $A_g(n)$ can be dropped. Indeed, let $d = m_1m_2 \mid n$ with $P^+(m_1) < n^\nu \le P^-(m_2)$ and $g(v(m_2; n)) \neq 0$, so that $v(m_2; n) \in G_1$. Since $m_1m_2 \mid n$, the vector $(v(m_2; n), v(m_1; n))$ is a subvector of a rearrangement of $v(n) \in C(R)$, and $C(R)$ is closed under rearrangement, because the parts of a coagulation may be listed in any order; the coordinates of $v(m_1; n)$ are less than ν. So the definitions of ψ and G_1 in [5, (7.1)] give $|v(m_1; n)| \le \psi(v(m_2; n)) \le \gamma - |v(m_2; n)|$, that is, $d \le n^\gamma$. The divisors of n are the products m_1m_2 of a divisor m_1 of n with $P^+(m_1) < n^\nu$ and a divisor m_2 of n with $P^-(m_2) \ge n^\nu$, each obtained once, so $A_g(n) = \sum_{m_1}\mu(m_1)\sum_{m_2} g(v(m_2; n))$. The first sum is $\mathbf{1}_{P^-(n)\ge n^\nu}$, and when it is 1 every divisor of n is an m_2. Hence

$$(*) \qquad A_g(n) = \mathbf{1}_{P^-(n)\ge n^\nu}\sum_{m\mid n} g(v(m; n)) \qquad (n \in \mathcal{N}),$$

the sum running over the ordinary divisors m of n. In particular, $A_g(n) = 0$ for $n \in \mathcal{N}$ with $P^-(n) < n^\nu$, and $A_g(n) = V_g(v(n))$ for squarefree $n \in \mathcal{N}$ with $P^-(n) \geq n^\nu$, whose divisors correspond to the subvectors of $v(n)$.

7.1. **A sieve bound with one parity term.**

Proposition 7.2. *Let $P = (1/2, 0, \nu)$ with $R(P)$ nonempty, and let g be admissible, extended by zero outside G_1, with $g(\emptyset) = 1$, and such that its restriction to ordered vectors is a finite sum of functions each supported on a convex polytope and Lipschitz continuous on it. For $c < 1$ put $\Phi(x) = V_g(x) + c\,(-1)^{\dim x}$. Let $\mathcal{N}_r = \{n \in \mathcal{N} : P^-(n) \geq n^\nu\}$. Assume, in addition to Ford and Maynard's hypotheses, parity balance on the primes and $\mathcal{N}_r$: $\left|\sum_{n \in \mathcal{P} \cup \mathcal{N}_r} \mu(n) w_n\right| \leq x(\log x)^{-B}$. Suppose that $\Phi \leq 0$ on $\mathcal{H}(P)$, and that $A_g(n) \leq 0$ on non-squarefree $n \in \mathcal{N}_r$ apart from an exceptional set E_x such that $\sum_{n \in E_x} |w_n| \leq 2x(\log x)^{-B}$ for every pair in Ford and Maynard's class at exponent B and every large x. Then $C^-(1/2, 0, \nu)$, taken over the class with the parity condition on $\mathcal{P} \cup \mathcal{N}_r$ adjoined, satisfies*

$$C^- \;\geq\; 1 + \frac{1}{1-c} \sum_{k \geq 2} \int_{\mathcal{H}(P) \cap \{x_1 \leq \cdots \leq x_k\}} \frac{\Phi(x)}{x_1 \cdots x_k}\, dx.$$

Proof. Put $F(n) = A_g(n) + c\,\mu(n) \mathbf{1}_{P^-(n) \geq n^\nu}$. A vector of $R(P)$ has components in $(0, 1/2)$ and sum 1, so it has at least three components, each a proper subsum and hence greater than ν; thus $\nu < 1/3$, and in particular $\nu \leq 1/2 = 1 - \gamma$, so Lemma 7.1 applies. Together with the parity hypothesis it gives $\sum_{\mathcal{P} \cup \mathcal{N}} w_n F(n) \ll_A x/\log^A x$ once $B \geq A$ is large enough. On primes $F = 1 - c$; on squarefree $n \in \mathcal{N}_r$, $F = \Phi(v(n)) \leq 0$, by $(*)$ and the hypothesis on Φ; on non-squarefree $n \in \mathcal{N}_r$, $F = A_g(n)$, which is ≤ 0 off E_x; and $F = 0$ on $\mathcal{N} \setminus \mathcal{N}_r$, where $A_g = 0$ by $(*)$. On $\mathcal{N}_r$ we have $\Omega(n) \leq 1/\nu$, so $|A_g(n)| \leq \|g\|_\infty 2^{1/\nu}$ by $(*)$, and F and $\Phi(v(n))$ are bounded there. Since $a_n \geq 0$, and the terms from E_x are $O(x(\log x)^{-B})$ by the hypothesis on E_x, $(1-c)\sum_p w_p \geq \sum_{n \in \mathcal{N}_r \setminus E_x} b_n F(n) - O_A(x/\log^A x)$.

The last sum is $(\sum_p b_p)$ times the sum of integrals in the statement, with error $O_g((\sum_p b_p)/B)$. Here $\mathcal{N}_r$ is the set of $n \sim x$ with $P^-(n) \geq n^\nu$ and $v(n) \in \mathcal{H}(P)$, and $\mathcal{H}(P) \cap \mathbb{R}^k$ is the set of x with $|x| = 1$ and every $x_i \in (\nu, 1-2\nu) \cup (2\nu, 1-\nu)$ [5, Lemma 8.4]. So its ordered vectors form a finite disjoint union of convex polytopes inside $C(R) \cap \{\nu \leq x_1 \leq \cdots \leq x_k\}$, where Ford and Maynard's condition (b.2) applies. For $J \subseteq [k]$ and one of the functions of the hypothesis on g, Lipschitz on a convex polytope $\mathcal{G}$, the ordered x of one of these polytopes with $x_J \in \mathcal{G}$ form a convex set, on which that function of x_J is Lipschitz. Applying (b.2) on each such set, as in the proof of [5, Lemma 7.20], and summing over J and the functions gives the part from $V_g(v(n)) = \sum_{J \subseteq [k]} g(v(n)_J)$, with error $O_g((\sum_p b_p)/B)$; (b.2) with $f = 1$ on the same polytopes gives the part from $c\,(-1)^k$. The faces with repeated coordinates are convex sets of projection measure zero, so (b.2) with $f = 1$, and $b_n \geq 0$, bound their b-mass by $O_g((\sum_p b_p)/B)$. As F and $\Phi(v(n))$ are bounded on $\mathcal{N}_r$, replacing $F(n)$ by $\Phi(v(n))$ at the non-squarefree $n \in \mathcal{N}_r$, including those in E_x, changes the sum by $O_g((\sum_p b_p)/B)$. Finally, by (b.1) the error $O_A(x/\log^A x)$ is $o(\sum_p b_p)$ for $A > \varpi$; then use Ford and Maynard's quantifiers with B arbitrarily large, as in the proof of [5, Theorem 7.3]. □

Lemma 7.3. *Let $P = (1/2, 0, \nu)$ with $0 < \nu < 1/3$. If B is large enough in terms of A, ϖ and ν, then (w), (I) and (II) imply $\sum_{n \notin \mathcal{P} \cup \mathcal{N}} \mu(n) w_n \ll x(\log x)^{-A}$. Hence global parity balance* (Par) *at a large enough logarithmic exponent implies parity balance on $\mathcal{P} \cup \mathcal{N}$ at any prescribed exponent.*

Proof. Only squarefree integers contribute to the sum. First consider those with $P^-(n) \leq n^\nu$, retaining the boundary because the components of $C(R)$ are strictly greater than ν. Writing $n = pm$ with p the least prime factor, their contribution is

$$-\sum_{\substack{pm \sim x,\ p \leq x^\nu \text{ prime} \\ P^-(m) > p,\ p^{(1-\nu)/\nu} \leq m}} \mu(m) w_{pm}.$$

Apply [5, Lemma 7.10], with designated block $\{m\}$ (on the sum $m > x^{1-\nu}/2 \geq x^{\nu/3}$), to the monomial inequality $p^{(1-\nu)/\nu}m^{-1} \leq 1$ and [5, Lemma 7.9] to $P^-(m) > p$, using the gap between an integer and $p + 1/2$ (as $m > 1$ on the sum, the value $P^-(1) = \infty$ does not enter). The resulting separated sums are covered by (II). Both parts of (w) supply the error bounds in those lemmas; the bounded losses of logarithmic powers are absorbed by choosing B sufficiently large. The other squarefree composite n have $P^-(n) > n^\nu$, so every component of $v(n)$ lies in $(\nu, 1-\nu)$. For $\nu \leq 1/4$ every such n lies in $\mathcal{N}$ by [5, Lemma 8.4] (at $\nu = 1/4$ a component $1/2$ cannot occur in a squarefree composite), so this exhausts the contributing complement of $\mathcal{P} \cup \mathcal{N}$.

If $1/4 < \nu < 1/3$, such an n has two or three prime factors, since $4\nu > 1$. With three, each component is less than $1 - 2\nu$, as the other two exceed ν, so $n \in \mathcal{N}$; with two, $v(n) = (\beta, 1-\beta)$ with $\beta < 1/2 < 2\nu$, which lies in $\mathcal{H}(P)$ if and only if $\beta < 1 - 2\nu$ [5, Lemma 8.4]. So the remaining integers are semiprimes qr, with $q < r$ and $q \geq (qr)^{1-2\nu}$. Let q run over primes in $[(x/2)^{1-2\nu}, \sqrt{x}]$. For each q, the conditions $qm \sim x$, $m > q$ and $m \leq q^{2\nu/(1-2\nu)}$ are interval conditions, so (I) bounds the sum over all such m. Subtract its composite cofactors, written uniquely as $m = pr$ with $p = P^-(m)$, $r \geq 2$ and $P^-(r) \geq p$; here $p \leq m^{1/2} \leq (qm)^\nu \leq x^\nu$ (as $r \geq p$), since $m \leq q^{2\nu/(1-2\nu)}$ says $m \leq (qm)^{2\nu}$, and we keep $p \leq x^\nu$ as a condition on p alone. [5, Lemmas 7.10–7.11], with designated block $\{q\}$ ($q \geq (x/2)^{1-2\nu} > x^{\nu/3}$; in the second condition the exponent $-2\nu/(1-2\nu)$ of q is at most -1, since $\nu > 1/4$), separate $q < pr$ and $pr \leq q^{2\nu/(1-2\nu)}$, and [5, Lemma 7.9] separates $P^-(r) \geq p$, using $p - 1/2$ (here $r \geq 2$, and the sum over q is kept in the coefficient). After grouping $e = qr$, these are (II) sums in p, e: since every such q exceeds $x^{1/3}$ for sufficiently large x, each $e \leq x$ has at most two possible prime divisors q, so the grouped coefficients are bounded. The resulting logarithmic and power-saving errors are $O_A(x/\log^A x)$ for large enough B. Subtraction leaves exactly the required semiprime sum. □

7.2. **Proof of Theorem B(a).** Take $c = -1/2$, and let $g = h\,\mathbf{1}_{G_1}$, where $h(\emptyset) = 1$, $h(y) = -1/2$ for $\dim y \in \{1, 2\}$ with $|y| \leq 1/2$, and $h = 0$ otherwise. Every component of a vector of $C(R)$ exceeds ν (see below), so $\psi \equiv 0$ and G_1 consists of the subvectors of vectors of $C(R)$ with sum at most $1/2$. For $\nu < 1/4$ this gives $G_1 \cap \mathbb{R}^1 = \{(y) : \nu < y \leq 1/2\}$, so that the single-component weight is $-1/2$ on $(\nu, 1/2]$, closed at $1/2$, and $G_1 \cap \mathbb{R}^2 = \{(y_1, y_2) : y_1, y_2 > \nu,\ y_1 + y_2 \leq 1/2\}$: for $\nu < y < 1/2$ and for such (y_1, y_2), the vectors $(y, (1-y)/2, (1-y)/2)$, $(1/4, 1/4, 1/4, 1/4)$, which has the coagulation $(1/2, 1/2)$, and (y_1, y_2, s, s) with $s = (1 - y_1 - y_2)/2$ lie in R, since their components lie in $(\nu, 1/2)$ and sum to 1, so that no proper subsum lies in $[0, \nu]$. For $1/4 \leq \nu < 1/3$, the composite vectors of $C(R)$ are the triples with all components in $(\nu, 1-2\nu)$ and the pairs $(\beta, 1-\beta)$ with $\nu < \beta < 1 - 2\nu$ [5, Lemma 8.4]; a two-component subvector of such a triple has sum greater than $2\nu \geq 1/2$, and the larger component of such a pair exceeds $2\nu \geq 1/2$, so $G_1 = \{\emptyset\} \cup \{(y) : \nu < y < 1 - 2\nu\}$. In each case g, on ordered vectors, is constant on finitely many convex polytopes, so it is admissible and satisfies the Lipschitz hypothesis of Proposition 7.2. Then $\Phi = 1 - N_1/2 - N_2/2 - (-1)^k/2$ on a k-dimensional vector, where N_1 and N_2 count its subvectors of dimension 1 and 2 with sum at most $1/2$, and $N_1 \geq k - 1$. Hence:

vector	Φ
$k = 2$, unequal components	0
$E = (1/2, 1/2)$, in $\mathcal{H}$ when $\nu < 1/4$	$-1/2$
$k = 3$, no component equal to $1/2$	0
$k = 3$, one component equal to $1/2$	$-1/2$
$k \geq 4$	$\leq 1 - 3/2 + 1/2 = 0$

So $\Phi \leq 0$ on $\mathcal{H}(P)$. On non-squarefree $n \in \mathcal{N}_r$ with $\nu < 1/4$, $A_g(n) = 1 - (D_1(n) + D_2(n))/2$, where D_1 counts the distinct prime divisors up to $\sqrt{n}$ and D_2 the divisors $d \leq \sqrt{n}$ with $\Omega(d) = 2$, squares included; this is ≤ 0 except at prime squares and cubes, where it equals $1/2$. Indeed, if $\Omega(n) \geq 4$, the two least prime factors $p \leq p'$ of n, counted with multiplicity, satisfy $p \leq n^{1/4}$ and $pp' \leq n^{1/2}$,

so $D_1(n), D_2(n) \geq 1$; if $n = p^2q$ with primes $p \neq q$, then $D_1(n) = 2$ unless $p \leq n^{1/4}$, and in that case p and p^2 are counted. Prime squares and cubes form the exceptional set: (I), applied with the modulus p and the one-point cofactor interval $\{p^{r-1}\}$, gives $\sum_{x/2<p^r\leq x} |w_{p^r}| \leq x(\log x)^{-B}$ for $r = 2, 3$. (For $\nu \geq 1/4$, the only non-squarefree members of $\mathcal{N}$ are p^2q with $p \neq q$, where $A_g = 0$, and p^3, where $A_g = 1/2$; see the description of $C(R)$ and G_1 above, and the proof of part (b) below for these values of A_g, which also hold at $\nu = 1/4$.)

In dimensions $4, 5, 6, 7, 8$, Φ is bounded below by -3, -6, -10, $-25/2$ and $-35/2$. In dimensions 5 to 8 this holds everywhere, since $N_1 \leq k$ and $N_2 \leq \binom{k}{2}$; in dimension 4 it holds off the hyperplanes $x_i + x_j = 1/2$, which have projection measure zero (on them it can fail: $\Phi(1/4, 1/4, 1/4, 1/4) = -9/2$). In dimensions 2 and 3, Φ vanishes almost everywhere. Every vector of $\mathcal{H}(P) \cap \mathbb{R}^k$ has its components in $(\nu, 1-\nu)$ [5, Lemma 8.4], and integration over ordered vectors is $1/k!$ times integration over all orderings, so the integral of $1/(x_1 \cdots x_k)$ over the ordered vectors of $\mathcal{H}(P) \cap \mathbb{R}^k$ is at most $I_k(\nu)$. For $P = (1/2, 0, \nu)$, $\mathcal{N}_r = \mathcal{N}$: a vector of $R(1/2, 0, \nu)$ has no proper subsum in $[0, \nu]$, so every component of a vector of $C(R)$ with at least two components exceeds ν, and hence $P^-(n) > n^\nu$ for every $n \in \mathcal{N}$. Ford and Maynard's class shrinks as B grows. Given ϖ and a constant below the right-hand side of Proposition 7.2, let B' be an exponent that the proposition supplies, and take $B \geq B' + 1$ so large in terms of B', ϖ and ν that Lemma 7.3 applies with $A = B' + 1$. Since $\mathcal{N}_r = \mathcal{N}$, a pair of the class at exponent B that satisfies (Par) then has $\left|\sum_{n\in\mathcal{P}\cup\mathcal{N}_r} \mu(n)w_n\right| \leq x(\log x)^{-B} + O(x(\log x)^{-B'-1}) \leq x(\log x)^{-B'}$ for large x, so it lies in the class of Proposition 7.2 at exponent B'; the bound on E_x holds at every exponent, by (I) as above. Hence, for $1/9 \leq \nu < 1/3$,

$$C^-(1/2, 0, \nu \mid \text{Par}) \geq 1 - 2I_4(\nu) - 4I_5(\nu) - \tfrac{20}{3}I_6(\nu) - \tfrac{25}{3}I_7(\nu) - \tfrac{35}{3}I_8(\nu),$$

with I_k as in Theorem B ($I_k = 0$ once $k\nu \geq 1$, so dimensions 9 and up do not occur, and $I_8 = 0$ for $\nu \geq 1/8$).

7.3. **The numbers.** Each $I_k(\nu)$ is non-increasing in ν: as ν increases the domain of integration shrinks, and the integrand is positive and does not depend on ν. A separately written program, from the audit of this theorem, encloses the $I_k(\nu)$ in exact integer arithmetic: $[0, 1]$ is cut into N cells of length $1/N$, the density $\mathbf{1}_{\nu<x<1-\nu}/x$ is bounded almost everywhere on each cell from below and above by integers at scale 10^{18}, the convolutions are formed as exact integer polynomial products and divided by k with outward rounding, the results are enclosed cell by cell with the support below $k\nu$ removed, and $I_k(\nu)$ is enclosed at 1. The program and its outputs are in the supplement's `thmB/` directory. With $N = 16000$ at $\nu = 3/20$ and $N = 9800$ at $\nu = 1/7$ this gives

$$C^-(1/2, 0, 3/20 \mid \text{Par}) \geq \tfrac{132606075869165383}{187500000000000000} > 0.7072324046,$$
$$C^-(1/2, 0, 1/7 \mid \text{Par}) \geq \tfrac{232429204005956881}{375000000000000000} > 0.6198112106,$$

and with $N = 8000$ at $\nu = 1/8$ it gives the lower bound $\frac{144641353317676797}{500000000000000000} > 0.2892827066$ there, with $I_7(1/8) \leq 9.15 \cdot 10^{-7}$. Below $\nu = 1/8$ the term I_8 enters. A copy of the program that adds it to the bound, with its enclosure arithmetic unchanged, gives with $N = 32000$ at $\nu = 0.116$ the lower bound $\frac{26309423794885631}{750000000000000000} > 0.0350792317$ there, with $I_8(0.116) \leq 8.49 \cdot 10^{-10}$; with the term I_8 set to zero, the copy reproduces the stored exact outputs at 3/20, 1/7 and 1/8 byte for byte. By monotonicity the bound is positive on the whole range $0.116 \leq \nu < 1/3$, at least 0.0350 throughout, at least 0.2892 for $\nu \geq 1/8$, at least 0.6198 for $\nu \geq 1/7$ and at least 0.7072 for $\nu \geq 3/20$. A floating-point integral recurrence, computed in the same audit, gives the underlying values 0.70749 at 3/20 and 0.62035 at 1/7; the same recurrence with the term I_8, run with the copy after the audit, gives 0.03551 at 0.116. This proves Theorem B(a). □

7.4. **Proof of Theorem B(b).** Let $1/4 < \nu < 1/3$. As in the proof of part (a), the composite vectors of $C(R)$ are the triples with all components in $(\nu, 1-2\nu)$, and the pairs $(\beta, 1-\beta)$ with

$\nu < \beta < 1-2\nu$ [5, Lemma 8.4]; moreover $\psi \equiv 0$ and $G_1 = \{\emptyset\} \cup \{(y) : \nu < y < 1-2\nu\}$. Take $g(\emptyset) = 1$, $g((y)) = -1/2$ for the one-component vectors of G_1, and $c = -1/2$. On squarefree $n \in \mathcal{P} \cup \mathcal{N}$ we have $A_g(n) - \mu(n)/2 = 1 - N_1(n)/2 - \mu(n)/2$, where $N_1(n)$ counts the prime factors p of n with $(\log p/\log n) \in G_1$; for a prime, a pair and a triple, $(N_1, \mu) = (0,-1), (1,1), (3,-1)$, so this is $\frac{3}{2}\mathbf{1}_{\mathcal{P}}(n)$. The only non-squarefree members of $\mathcal{N}$ are p^2q with $p \neq q$ and p^3. Their ordinary divisors m with $v(m;n) \in G_1$ are $1, p, q$ and $1, p$ respectively, so $A_g(p^2q) = 0$ and $A_g(p^3) = 1/2$, and $\mu = 0$ at both. Lemma 7.1, together with (Par) and Lemma 7.3, gives $\sum_{n \in \mathcal{P}\cup\mathcal{N}} w_n(A_g(n) - \mu(n)/2) \ll_A x(\log x)^{-A}$, that is,

$$\tfrac{3}{2} \sum_{x/2 < p \le x} w_p + \tfrac{1}{2} \sum_{x/2 < p^3 \le x} w_{p^3} \ll_A x(\log x)^{-A}.$$

The sum over cubes is $O(x(\log x)^{-B})$ by (I), applied with intervals that isolate the cofactor p^2, as in part (a). Hence $\sum_{x/2<p\le x}(a_p - b_p) \ll_A x(\log x)^{-A}$ once B is large enough, which is the first assertion of part (b). Since $\sum_p b_p \ge x(\log x)^{-\varpi}$ by (b.1), taking $A > \varpi$ gives $C^-(1/2, 0, \nu \mid \mathrm{Par}) \ge 1$. The pair $a_n = b_n = 1$ on $(x/2, x]$ has $w = 0$, so it satisfies (w), (I), (II) and (Par), and (b.1) and (b.2) hold for it by [5, Lemma 4.6] with $q = 1$, $y = x/2$ (§2); as $\sum_p a_p = \sum_p b_p$ for it, $C^-(1/2, 0, \nu \mid \mathrm{Par}) \le 1$. □

7.5. **The remaining range.** Let $1/3 \le \nu \le 1$. We show directly that (w), (I) and (II), with $\gamma = 1/2$ and $\theta = 0$, imply $\sum_{x/2<p\le x} w_p \ll_A x(\log x)^{-A}$ when B is sufficiently large. For $x/2 \le t \le x$, put

$$\Pi(t) = \sum_{x/2 < p \le t} w_p \frac{\log p}{\log x}.$$

Apply [5, Lemma 7.12] with $k = 1$ (its $\ell = 6\lceil 1/(1-\gamma)\rceil$ is 12, as $\gamma = 1/2$), with $\psi_{u,v} = \mathbf{1}_{u=v=1}$ and $\mathcal{T}_{1,1} = [\log(x/2)/\log x, \log t/\log x]$; the other regions may be empty, and the sum of that lemma is then $\Pi(t)$. Its hypotheses are (w) and (I), with B sufficiently large, and the indicated region is convex. Its implied constant depends only on k and the fixed parameters, not on the region, so it gives, uniformly in t,

$$|\Pi(t)| \ll (\log x)^{12} \sup_{|\beta_i| \le 1} |S_{\boldsymbol{\beta}}(t)| + x(\log x)^{-B+1}, \qquad S_{\boldsymbol{\beta}}(t) = \sum_{\substack{x/2 < n = d_1 \cdots d_{12} \le t \\ d_i < n/\sqrt{x} \ (1 \le i \le 12)}} w_n \prod_{i=1}^{12} \beta_i(d_i).$$

Every tuple in this sum has at least three factors greater than 1. Indeed, the factors greater than 1 have product $n > 1$ and are each less than $n/\sqrt{x} < n$, so there are at least two of them; if there were exactly two, we would have $n < (n/\sqrt{x})^2 = n \cdot n/x \le n$, a contradiction. The least factor greater than 1 is therefore at most $n^{1/3} \le x^{\nu}$. Partition the tuples by the first index j for which $1 < d_j \le x^{\nu}$. For $i < j$, the conditions $d_i = 1$ or $d_i > x^{\nu}$ depend on one variable only and are absorbed into its coefficient.

For each j, remove the conditions $d_i < n/\sqrt{x}$ and $n \le t$ using [5, Lemma 7.9]. Write them as

$$\frac{n}{d_i} > \lfloor \sqrt{x} \rfloor + \frac{1}{2}, \qquad n < \lfloor t \rfloor + \frac{1}{2}.$$

The compared quantities are positive, are bounded by $2x$, and have gaps at least $1/2$. For the condition on d_i, the two variables of that lemma are n/d_i, the product of the other eleven factors, and d_i, with the threshold as a constant function; for the last condition one may group the variables by their total product and use a dummy second variable equal to 1. Constant unimodular factors of the twists are taken out of the absolute value. The grouped coefficient arrays have absolute value at most $|w_n|\tau(n)^C < x^2$ for a fixed C and large x, by (w) and the divisor bound. Thus all the hypotheses of [5, Lemma 7.9] hold. Its twists, $(n/d_i)^{it_i}$ and n^{it_0}, factor into one-variable

coefficients of unchanged absolute value. After these thirteen separations, at a cost $O((\log x)^{13})$ and a negligible error, the sums have the form

$$\sum_{\substack{x/2<mr\le x\\ 1<m\le x^{\nu}}} \xi(m)\kappa(r)w_{mr}, \qquad |\xi(m)|\le 1, \quad |\kappa(r)|\le \tau(r)^{12},$$

where $m = d_j$ and $\kappa(r)$ sums over the ordered factorisations of r into the other eleven factors, so that $|\kappa(r)|$ is at most their number $\tau_{11}(r) \le \tau(r)^{10}$ [5, Lemma 7.7]. Condition (II) bounds each sum by $x(\log x)^{-B}$ for $B \ge 12$. Consequently $\sup_t |\Pi(t)| \ll_A x(\log x)^{-A}$ once $B \ge A + 26$ is also large enough for the preceding lemmas. Partial summation with the bounded function $\log x/\log t$, whose total variation on $[x/2, x]$ is bounded, gives $\sum_{x/2<p\le x} w_p \ll_A x(\log x)^{-A}$. Taking $A > \varpi$ and using (b.1), we get $C^-(1/2, 0, \nu) \ge 1$ for $1/3 \le \nu \le 1$, and adjoining (H_α) can only increase C^- (§2). For $0.116 \le \nu < 1/3$, part (a) applies, since (H_α) implies (Par). This proves the final assertion of Theorem B. □

8. Limits of the construction

The certified cases stop at 4/15 for a computational reason. Past 1/3 the family (1) meets a structural obstruction, and we conjecture that it cannot pass 1/3; how close to 1/3 it can get is not known. In floating-point experiments the obstruction appears to belong to the family rather than to the problem: on three of the four grids we could compute, a discretised version of the problem with fully free data (Type I rows up to and including level 1/2, values at most 4 times the baseline; see below) gets past 1/3, while for richer data that see the largest small component and $\eta \le 1/36$ the values at the least level computed at or past 1/3 on each grid lie within 0.05 of -1 (below it on two of the four grids under both boundary conventions tested, with the other two unresolved under the tested conventions and boxes), and at the next grid point between -0.96 and -0.78. We keep the proved, measured and conjectural parts apart.

The next case. An uncertified numerical candidate occurs at the natural next case, $(\eta, \alpha') = (1/36, 17/60)$, which has $m^* = 10$ forced levels: with the profile $w_2 = 1$, $w_3 = -2$, constants designed on a grid of step 1/1800 and checked on grids of step 1/1440 and 1/2160, after re-forcing the functions and adjusting two class constants to impose $T_0 = 0$, give a positive numerical margin $1 - \min f/f(1)$ of a few hundredths (between 0.03 and 0.05 at the two resolutions; floating point, from one design code, not independently recomputed; the supplement holds the constants but not the design code or its evaluation logs). We do not claim this case. With the second-order engine of §5 as it stands, a certificate is out of reach at affordable cost, for two reasons. First, affected one-step couplings (the terms $C_{j,1}$ of (2)) have a first-order enclosure loss when a threshold window straddles the jump of the rough kernel at α'; at fixed block width this loss is not removed by increasing M, and increasing the number of blocks reduces it in the measured runs. A control run at $(1/30, 4/15)$ shows no comparable loss. The likely reason, which we have not proved, is that such cells occur only when $2\alpha' - 1/2 > \eta$, while at $(1/30, 4/15)$, $2\alpha' - 1/2 = \eta$ exactly. Second, the widths produced by the second-order engine compound down the forcing chain, by a measured factor of about 3 per level from level 6 to level 1: at $M = 3600$ the width of h_1 is about 3.6, against a margin of a few hundredths. Between $M = 3600$ and 5400 the widths at levels 3 to 5 shrank approximately as $M^{-1.5}$; assuming a quadratic rate would be optimistic, and we do not supply a cost projection. The measurements suggest treating the jump exactly and reducing the propagated enclosure widths; they do not establish that these changes are necessary or sufficient. The certified cases currently end at 4/15.

The obstruction past 1/3**.** The construction of §3 adds to Ford and Maynard's class constants one forced function h_j for each number j of small components, a function of their total size. We

conjecture that such constructions cannot pass twisted level 1/3 (Conjecture 8.2 below). The proved part is the following.

Lemma 8.1 (Rows that the forced functions cannot reach)**.** *Let $1/3 \le \alpha' < 1/2$, and let s be a tuple of small components with total $\sigma \in [1-2\alpha', \alpha']$, an interval that is nonempty exactly when $\alpha' \ge 1/3$.*

(a) *Every free vector that contains the components of s has at most one rough component, and its small total exceeds $1/2$.*

(b) *Let g be a bounded function on free vectors that vanishes on those with small total greater than $1/2$. Then the extension of g by (6.3) vanishes at every vector that contains the components of s, and so $T_t(g)(s) = 0$.*

In particular, for free data (1) the functions h_j do not enter the twisted rows $T_t(f)(s)$ at these tuples: only the class constants do.

Proof. (a) Let u be a free vector that contains the components of s, and let z consist of its other components. They total $1-\sigma \in [1-\alpha', 2\alpha']$, and each lies in $[\eta, 1/2)$, so each is small or rough. Two rough components of z would total more than $2\alpha' \ge 1-\sigma$, so z has at most one. If z had no small component, it would be a single rough component of size $1-\sigma \ge 1-\alpha' > 1/2$, which is impossible. So the small total of u is 1 if z has no rough component, and $1-r > 1/2$ if it has one, r.

(b) By (a), g vanishes at every free vector that contains the components of s. A vector that contains them and is not free is not E, whose components are $1/2 > \alpha'$, so it has exactly one large component a. Relation (6.3) expresses its value through the free values at the vectors obtained by fragmenting a into components below $1/2$. Each of these is a free vector that still contains the components of s, so the extension of g vanishes there as well. The functional $T_t(\cdot)(s)$ integrates the values at the vectors (s, z), all of which contain the components of s, so $T_t(g)(s) = 0$. For the last statement, apply (b) to the forced part $\mathbf{1}[1 \le j \le m^*,\ \tau \le \alpha']\, w_k h_j(\tau)$ of (1), which vanishes when $\tau > \alpha'$, and use the linearity of the extension and of T_t. □

The class constants enter these rows through smoothing kernels, so for the free data (1) the rows of Lemma 8.1 are linear conditions on the finitely many constants alone, one for each σ in an interval. (In the grid programmes below, the rows at $\sigma = 1-2\alpha'$ also contain the added function of their own label, through the closure node with two rough components equal to α'; see the README.) The rest of this section reports floating-point experiments. Except for the rigorous certificates named below, every value in it is a floating-point solution of a finite linear programme: evidence, not proof. In the experiments on families, the programme minimises $f(1)$ over the class constants and the added functions, tabulated on a grid of step $1/N$, subject to the twisted rows and to $f \ge -1$ at every generated non-prime grid point, including the closure-labelled free nodes and the extended points described in the README, except at E read as a rough component plus a large one (there is no upper positivity condition). The problem itself restricts the data only through $f \ge -1$ and boundedness, so each programme carries an explicit box, which we state and vary: the class constants lie in $[-1, U]$ and the added functions are bounded by C in absolute value. For $\alpha' \ge 1/3$, allowing the constants of the classes whose free vectors all carry added data to go down to $-L < -1$ adds nothing to the family itself. Adding a constant δ_j to $c[j,2,b]$ for every bin b and subtracting δ_j from the added function of level j leaves f unchanged at every free vector, and the same holds for the classes with one rough component. So that box gives no more than the box with $U+L-1$ and $C+L-1$ in place of U and C. On the grid this fails at least at the nodes where the one rough component equals $1/2$ and at the G_j columns that the README lists as omitted, which carry quadrature weight but no added variable; for $c[j,2,b]$ the exchange is exact on the grids we checked ($\alpha' = 24/72$ and $25/72$). We therefore bound every constant below by -1. Unless stated otherwise, the family programmes below use $U = 4$, with $C = 50$ for the richer families. All the

families below have free data of the form

$$f(u) = c[j,k,b] + \mathbf{1}[\tau \le \alpha']\, H_j(\tau, \mathfrak{m}) + \mathbf{1}[k=1]\, G_j(\tau, \mathfrak{m}), \tag{4}$$

with j, k, b and τ as in (1) and $\mathfrak{m}$ the largest small component: the forced family (1) with $w_k = 1$ is the case where H_j does not depend on $\mathfrak{m}$ and $G_j = 0$, and the family with G_j the case where neither depends on $\mathfrak{m}$. The values are tabulated at the grid points i/N, and every integral in the rows and in (6.3) is a composite trapezoid sum over the grid points of its closed range, so that a component equal to α' counts half as small and half as rough, and a component equal to $1/2$ half as rough and half as large. The kernels of (6.3) are evaluated at the endpoint $a = 1/2$ of the large component by their limits from above, for the class constants and the added functions alike, so that each programme depends on the data only through the values of f at the grid points. The remaining grid conventions of these programmes (closure points, the vector E, and the added functions they leave out) are listed in the supplement's experiments README, under "Conventions of the family programmes". Coefficients of absolute value below 10^{-9} are deleted from the matrices; on the resulting matrices the rows of Lemma 8.1 have right-hand-side tolerance 10^{-9}, and the other twisted rows are equalities. Accepted solves have status Optimal, equality residuals, positivity violations and box violations at most 10^{-7}, and residuals of the rows of Lemma 8.1 at most $10^{-9} + 10^{-10}$. A value below -1 is quoted as a numerical prime-free candidate only if it passes these checks; on the matrices before the deletion, the residuals of the retained twisted rows and the generated positivity violations are at most $5 \cdot 10^{-6}$ in every retained-row case quoted. In the explicitly stated comparison without the rows of Lemma 8.1, those rows are exempt from the residual requirement. In every solution quoted, $f \ge -1 - 10^{-9}$ at every free grid point, the end points of the intervals included. Each value comes with a binary64 weak-duality lower bound for its programme, computed on the matrices before the deletion. A value above -1 whose bound is also above -1 is numerical evidence that the programme has no candidate. A separate check in an earlier discretisation of the forced family, on grids of step $1/360$ and $1/720$, agrees with the lemma: at the points checked, the coefficients of the forced functions in the rows of Lemma 8.1 are at most $4 \cdot 10^{-16}$, which is binary64 rounding, except within three grid steps of the endpoint $\sigma = 1 - 2\alpha'$, where the discretisation leaves errors of the order of the grid step.

For the forced family (1) with the profile $w_k = 1$ at $\eta = 1/36$, on a grid of step $1/72$, with two bins for the rough component (split at $30/72$ for $\alpha' \le 1/3$ and at $31/72$ past $1/3$) and the forced functions bounded by 10^4, the best $f(1)$ is -0.910, -0.778, -0.168 and -0.004 at $\alpha' = 23/72$, $1/3$, $25/72$ and $13/36$ for $U = 4$; for $U = 1000$ each value changes by less than $2 \cdot 10^{-4}$. Without the rows of Lemma 8.1 the values at $25/72$ and $13/36$ are -0.949 and -0.964 ($U = 4$). On this grid, retaining those rows raises the objective sharply at the two tested points past $1/3$; this does not prove a bound for the continuum family. With bins of one grid step the values at $25/72$ and $13/36$ are -0.258 and -0.186 for $U = 4$, and -0.265 and -0.187 for $U = 1000$: the collapse persists, but its size depends on the bins. On a grid these values depend on the conventions at the boundaries, which are load-bearing there: the twisted rows include $\sigma = \alpha'$; Type I enters only through (6.3), so there is no condition at a modulus of size exactly $x^{1/2}$; the positivity row at E read as a rough component plus a large one is dropped; and a component equal to α' or $1/2$ is weighted half on each side. E read as two rough components is the only vector of its class, whose constant enters only T_0, with coefficient $1/(2N)$. That constant is bounded below by -1 like every other, and the bound is active at every optimum reported in this section, so the parity row reads $T_0' x = 1/(2N)$, where T_0' is T_0 without that term. Fixing the constant at 0, so that $T_0' x = 0$, raises the values by 0.0029 to 0.0033 at the five points we tested ($25/72$ for the forced family, for the family with G_j at two boxes and for the richer family, and $1/3$ for the richer family; see the README). These boundary conventions can materially affect the finite programmes. The runs reported here do not bound the change produced by another convention, or the error in locating a continuum obstruction. Under the stated conventions, the two tested forced-family values past $1/3$ remain far above -1.

Conjecture 8.2. *Let $0 < \eta < \alpha' < 1/2$, fix finitely many bins, and let f be the extension by (6.3) of free data of the form (1), with any profile (w_k), any m^*, any class constants and any bounded measurable functions h_j on $[j\eta, \alpha']$. Then there is no $\alpha \in (1/3, \alpha']$ such that f satisfies (3) for some $m > 0$, $T_0(f) = 0$, and $T_t(f)(s) = 0$ for every $t \geq 1$ and every tuple s with $s_i \geq \eta$ and $|s| \leq \alpha$.*

Theorem 4.1 uses (3) only at vectors with no subsum equal to $1/2$; the conjecture is stated for the full form, which is the form the experiments test.

The conjecture rests on Lemma 8.1 and on the experiments above, which cover $\eta = 1/36$, one grid, $\alpha = \alpha'$ at two points, and the bins stated. The profile is no restriction there: for $\alpha' \geq 1/3$ a free vector with small total at most α' has exactly two rough components and small total below $1 - 2\alpha'$, so only w_2 and the values of h_j on $[j\eta, 1 - 2\alpha')$ enter. The lemma applies to any bounded added data that vanish on free vectors with small total above $1/2$, not only to (1). It does not apply to data that vary with the small total on vectors with exactly one rough component r, where $\tau = 1 - r > 1/2$: completing a tuple of Lemma 8.1 by one rough component r leaves small mass $1 - \sigma - r > 0$, so such data can enter those rows when the remaining mass admits a completion by small components, and then only through smoothing kernels. The class constants $c[j, 1, b]$ are data of this kind, constant on each bin of $r = 1 - \tau$. Since Conjecture 8.2 allows any finite partition into bins, it covers such data when they are step functions of τ with finitely many steps, and no others.

Two richer families of this kind were tried, at $\eta = 1/36$ on a grid of step $1/72$ unless stated otherwise. Adding free functions $G_j(\tau)$ on vectors with exactly one rough component, which still see the small primes only through their number and total size, gives -0.272 at $\alpha' = 25/72$ for $U = 4$, $C = 50$, and -0.410 for $U = C = 50$, 200, 1000 and 5000 (these programmes omit G_j at small total exactly $1/2$ and some additional feasible high-j columns, as specified in the README; these are restrictions, and the value also depends on the treatment of E). As functions of τ these contain the finite step functions that Conjecture 8.2 covers with finer bins; the programme treats G_j as values at grid points, a different discretisation from the forced family with bins of one grid step reported above, which gives -0.265 at the same point for $U = 1000$, 0.15 higher. For such data the rows of Lemma 8.1 are an integral equation of the first kind, with smoothing kernels; that is consistent with values well above -1, but proves nothing. Free data that also see the largest small component $\mathfrak{m}$, through functions $H_j(\tau, \mathfrak{m})$ for $\tau \leq \alpha'$ and $G_j(\tau, \mathfrak{m})$ on vectors with one rough component (on some grids with omitted G_j columns as above; see the README), were computed at the ten points of Table 3, for $U = 4$ and $C = 50$, and on the grid of step $1/72$ also for $U = C = 1000$. For $\eta \leq 1/36$, at the least level computed at or past $1/3$ on each grid ($1/3$ for $\eta = 1/36$, $31/90$ for $\eta = 1/45$, $17/48$ for $\eta = 1/48$; no run was made at $1/3$ for the last two), the values lie within 0.05 of -1: below it, with either treatment of E, on the grid of step $1/108$ at $\eta = 1/36$ and at $\eta = 1/45$, and within 0.003 of it on the other two grids. At $\eta = 1/48$ the value changes side when E's constant is fixed at 0. On the grid of step $1/72$ at $\eta = 1/36$, for $U = 4$, $C = 50$ both displayed values exceed -1, but the original weak-duality bound lies below -1; for $U = C = 1000$, fixing E's constant at 0 changes the value from below -1 to above it. At the next grid point the values lie between -0.96 and -0.78. By the acceptance rule above, the two values below -1 with either treatment of E are numerical prime-free candidates of this family on their grids. We draw no conclusion from them.

In a discretised version of the whole problem, every factorisation shape on a grid of step $1/K$ (a partition of K into parts of at least ηK) carries its own value, subject to the Type I rows at level $1/2$, with the level closed ($|v| \leq K/2$, i.e. Type I at moduli of size exactly $x^{1/2}$; on a grid these rows include the one that couples the vector E to the other shapes with a part $K/2$, which the family programmes above do not carry), and the twisted rows up to level α, with values between 0 and 4 times the baseline. On three of the four grids we computed, this problem has prime-free solutions past $1/3$; on the fourth, $K = 30$, it has one at $1/3$ and none at the next grid point. At $\eta = 1/36$ it has one at $\alpha = 13/36$ and none at $14/36$ for $K = 36$, and one at $26/72 = 13/36$ and none at $27/72$ for $K = 72$; at $\eta = 1/30$ it has one at $1/3$ and none at $11/30$ for $K = 30$, and one at

η	grid step	α'	best $f(1)$	weak-duality bound	with E's constant at 0
1/36	1/72	1/3	$-0.9999986^{\dagger}$	-1.0000349	-0.9970
1/36	1/72	25/72	-0.9025	-0.9027	
1/36	1/108	1/3	-1.0139	-1.0141	-1.0119
1/36	1/108	37/108	-0.9483	-0.9484	
1/45	1/90	31/90	-1.0415	-1.0417	-1.0391
1/45	1/90	32/90	-0.9578	-0.9580	-0.9555
1/48	1/96	17/48	$-1.0004^{\dagger}$	-1.0006	-0.9981
1/48	1/96	35/96	-0.7828	-0.7831	
1/30	1/90	1/3	-0.8847	-0.8848	
1/27	1/108	1/3	-0.7973	-0.7974	

TABLE 3. The family whose free data also see the largest small component, for $U = 4$, $C = 50$ (floating point; the last column fixes the constant of E read as two rough components at 0, so that $T_0'x = 0$; values are rounded up and bounds down). †On the grid of step 1/96, the two treatments of E give values on opposite sides of -1. On the grid of step 1/72, for $U = 4$, $C = 50$, both values are above -1, but the original weak-duality bound is below it. At the latter grid point with $U = C = 1000$, the value is -1.0000649, and -0.9971 with E's constant at 0.

21/60 and none at 22/60 for $K = 60$. Each "none" is stated for Type I at $|v| \le K/2$ (closed) and values at most 4 times the baseline; both are load-bearing. For $K = 36$, with the Type I level open ($|v| < K/2$) and values at most 4 times the baseline, the problem has a floating-point prime-free solution at 14/36 and a certified "none" at 15/36; for $K = 60$, values up to 100 times the baseline give one at 22/60 (below). These comparisons do not give a uniform grid-step error bound, and the other grids were not recomputed under either change. For $K = 36$ and $K = 72$ at $\eta = 1/36$ and $\alpha = 13/36$, prime-free solutions exist exactly: each of these two programmes has a rational solution whose value at the prime is exactly 0, certified by separately written checkers in exact integer and outward-rounded arithmetic, from data that the supplement's scripts regenerate. The other two prime-free solutions, at $\eta = 1/30$, are floating-point solutions with row residuals below $3 \cdot 10^{-15}$. Each "none" carries a rigorous certificate: a dual solution whose weak-duality bound over all shapes is evaluated in integer arithmetic, with a rigorous allowance for the final rounding, computed by the programme that found the dual solution; the dual vectors are not included in the supplement and the bound was not recomputed by separate code. These statements are for values at most 4 times the baseline: with values up to 100 times the baseline, at $\eta = 1/30$ and $K = 60$, the problem has a prime-free solution at 22/60 as well (floating point, row residuals below $3 \cdot 10^{-16}$, values up to about 99 times the baseline). At $(\eta, \alpha) = (1/36, 13/36)$, letting the values depend on the small parts only through their number and total, and on the other parts only through the number of rough parts, the largest rough part and the large part (a family that contains the grid analogue of (1)), raises the least weight on the primes to about 0.83 times the baseline for $K = 36$ and 0.81 for $K = 72$. Resolving the largest small part as well lowers these to 0.41 and 0.45, and resolving the two largest to 0.03 and 0.096; with the three largest the value is exactly zero for both $K = 36$ and $K = 72$, since the exact solutions above take one value per class of this family. Whether a bounded number of largest small parts suffices as the grid is refined is open. These grid results are consistent with Conjecture 8.2, and they suggest that the obstruction at 1/3 belongs to the families rather than to the problem; neither is shown.

Conjecture 8.2 concerns this family of constructions. It says nothing about whether Hooley's condition, with Type I information at level 1/2 and a small Type II range, forces primes at any level not covered by Theorem A; for instance, for $\nu < 1/30$ that question is open at every $\alpha \in [4/15, 1/2]$; above 1/2 we take no position (§1).

9. Consequences for Goldbach

Let N be a large even integer and $L = \log N$. The Goldbach count is $G(N) = \sum_{m<N} \Lambda(m)\Lambda(N-m)$, and the conjectured size is $G(N) \sim \mathfrak{S}(N)N$, where

$$\mathfrak{S}(N) = 2C_2 \prod_{p\mid N,\ p>2} \frac{p-1}{p-2}, \qquad C_2 = \prod_{p>2}\Big(1 - \frac{1}{(p-1)^2}\Big).$$

Read as a sum over m, $G(N)$ pairs the von Mangoldt function with the sequence $m \mapsto \Lambda(N-m)$. The constructions of the previous sections reflect to this setting. We replace $\Lambda(N-m)$ by a nonnegative sequence that agrees with a local model of $\Lambda(N-m)$ in a list of statistics, but vanishes on all large primes m. The listed statistics, with Hooley's twisted condition at the levels of Theorem A, therefore do not by themselves force a Goldbach-scale count.

9.1. **The local model.** Put $S = \mathfrak{S}(N)$, $w = \exp(\sqrt{L})$ and $P_w = \prod_{p\le w,\ p\nmid 2N} p$, and define, for $1 \le m < N$,

$$b'_m = K_w\, \mathbf{1}_{(m,N)=1}\, \mathbf{1}_{(N-m,P_w)=1} \prod_{p\mid m,\ p>w} \frac{p-1}{p-2}, \qquad K_w = S \prod_{p\le w,\ p\nmid 2N} \frac{p-1}{p-2} \sim e^{\gamma}\sqrt{L}.$$

This is a presieved Euler model of $\Lambda(N-m)$: its masses in intervals and in residue classes to small moduli are those of $\Lambda(N-m)$.

Lemma 9.1 (Class masses of b')**.** *Fix $A, C_0 > 0$. Uniformly for intervals $I \subseteq [1, N)$, integers $1 \le q \le L^{C_0}$ and residues r,*

$$\sum_{m\in I,\ m\equiv r \bmod q} b'_m = \frac{|I|}{\varphi(q)}\, \mathbf{1}_{(N-r,q)=1} + O_{A,C_0}(NL^{-A}).$$

For $\Lambda(N-m)$ in place of b'_m the same estimate is the Siegel–Walfisz theorem; in the classes with $(N-r,q) > 1$ only powers of primes dividing q occur, and they contribute $O(L^2)$.

We use upper and lower sieve weights of dimension 1 in the form of the fundamental lemma of sieve theory [13, Lemma 6.3], [7, §6.5]. Let $D = \exp(L^{3/4})$. There are real numbers $\lambda_e^{\pm}$, chosen once from D and w, with $\lambda_1^{\pm} = 1$, $|\lambda_e^{\pm}| \le 1$ and $\lambda_e^{\pm} = 0$ for $e \ge D$, such that $\sum_{e\mid n} \lambda_e^{-} \le \mathbf{1}_{n=1} \le \sum_{e\mid n} \lambda_e^{+}$ for every squarefree n whose prime factors are at most w, and such that, for large N, for every set $\mathcal{P}$ of primes $p \le w$ and every multiplicative function g with $0 \le g(p) < 1$ that satisfies the dimension condition $\prod_{y_1\le p<y_2,\ p\in\mathcal{P}}(1-g(p))^{-1} \le (\log y_2/\log y_1)(1 + K/\log y_1)$ for all $2 \le y_1 < y_2$,

$$\sum_{e\mid P} \lambda_e^{\pm}\, g(e) = \big(1 + O_K(e^{-L^{1/4}})\big) \prod_{p\in\mathcal{P}} \big(1-g(p)\big), \qquad P = \prod_{p\in\mathcal{P}} p, \tag{5}$$

since $\log D/\log w = L^{1/4}$. We use $g(p) = 1/p$, and $g(p) = 1/(p-1)$ on sets of odd primes; both satisfy the dimension condition with an absolute K. For a product P of primes $p \le w$ put $\theta_P^{\pm}(n) = \sum_{e\mid(n,P)} \lambda_e^{\pm}$, so that $\theta_P^{-}(n) \le \mathbf{1}_{(n,P)=1} \le \theta_P^{+}(n)$.

Proof of Lemma 9.1. Let $P(w) = \prod_{p\le w} p$. Every prime $p \le w$ with $p \mid 2N$ divides N, since N is even, and for such p the conditions $p \nmid m$ and $p \nmid N-m$ coincide. Hence $\mathbf{1}_{(m,N)=1}\mathbf{1}_{(N-m,P_w)=1} = \mathbf{1}_{(N-m,P(w))=1}$ unless m has a prime factor $p \mid N$ with $p > w$. At most $L/\log w = \sqrt{L}$ primes $p > w$ divide N, so at most $\sqrt{L}\,N/w$ integers $m < N$ are of this kind. Every $m < N$ has at most $\sqrt{L}$ prime factors $p > w$, so the last product in b'_m is $1 + O(\sqrt{L}/w)$. As $K_w \ll \sqrt{L}$,

$$\sum_{m\in I,\ m\equiv r \bmod q} b'_m = K_w M + O(NL/w), \qquad M = \#\{n \in N - I : n \equiv N - r \bmod q,\ (n, P(w)) = 1\}.$$

Every prime factor of q is at most $L^{C_0} < w$, so $M = 0$ if $(N-r,q) > 1$. Otherwise the condition $(n, P(w)) = 1$ only concerns the primes $p \le w$ with $p \nmid q$; let P be their product. By the Chinese

remainder theorem an interval of length $|I|$ contains $|I|/(qe)+O(1)$ integers $n\equiv N-r \pmod q$ with $e\mid n$, for every $e\mid P$. So M lies between the two numbers

$$\frac{|I|}{q}\sum_{e\mid P}\frac{\lambda_e^{\pm}}{e}+O(D)=\frac{|I|}{q}\prod_{\substack{p\le w\\ p\nmid q}}\Big(1-\frac{1}{p}\Big)\big(1+O(e^{-L^{1/4}})\big)+O(D),$$

by (5) with $g(p)=1/p$. The product is $\frac{q}{\varphi(q)}V(w)$, where $V(w)=\prod_{p\le w}(1-1/p)$. Since $2(1-\frac12)=1$ and $\big(1-\frac{1}{(p-1)^2}\big)\frac{p-1}{p-2}\big(1-\frac1p\big)=1$ for $p>2$, the factors at the primes $p\le w$ cancel, and

$$K_wV(w)=\prod_{p>w}\Big(1-\frac{1}{(p-1)^2}\Big)\prod_{p\mid N,\ p>w}\frac{p-1}{p-2}=1+O(\sqrt{L}/w).$$

Hence $K_wM=\frac{|I|}{\varphi(q)}\mathbf{1}_{(N-r,q)=1}+O(Ne^{-L^{1/4}}+\sqrt{L}\,e^{L^{3/4}})$, and every error term is $O_A(NL^{-A})$. □

If every prime factor of m exceeds some $y\ge w$, the last product in b'_m is $1+O(L/(y\log y))$. So on integers without small prime factors, b' is $K_w(1+O(L/(y\log y)))$ times the constant baseline $b=1$ of §2, times $\mathbf{1}_{(m,N)=1}$ and a presieve condition on $N-m$ with small moduli. The proof below transfers the estimates of §4 through that presieve.

9.2. **The reflected sequences.**

Proposition 9.2. *Let (η,α') be one of the constructions of Table 4 below. Fix $0\le\alpha<\alpha'$, $0<\nu<\eta$ and $C>0$. For every large even N there is a sequence $a_m\ge0$ on $[1,N)$, depending only on the construction, α, ν, C and N, that has the following properties for all $A,B>0$. Implied constants depend only on the construction and on α, ν, A, B and C.*

(G0) *$a_m=0$ for every prime $m\ge N/L^C$. Moreover, $\sum_m a_m\Lambda(m)\ll SNL^{-C}+N^{1/2}L$.*

(G1) *(Type I at the Bombieri–Vinogradov level, with interval maxima)*

$$\sum_{d\le N^{1/2}L^{-C/2}}\tau(d)^B\max_I\Big|\sum_{m\in I,\ d\mid m}(a_m-b'_m)\Big|\ll NL^{-A},$$

the maximum over intervals $I\subseteq[1,N)$.

(G2) *(Type II on small factors) For all real coefficients ξ_ℓ and κ_k,*

$$\sum_{1<\ell\le N^{\nu}}\sum_{k<N/\ell}\xi_\ell\,\kappa_k\,(a-b')_{\ell k}=0.$$

(G3) *(Hooley's condition at level α)*

$$\sum_{q\le N^{\alpha}}\Big|\sum_{m<N,\ q\mid m}\mu(m)(a_m-b'_m)\Big|\ll NL^{-A}.$$

(G4) *(Parity, on every interval) For $f\in\{\mu,\lambda\}$, uniformly in intervals $I\subseteq[1,N)$,*

$$\sum_{m\in I}(a_m-b'_m)f(m)\ll NL^{-A}.$$

(G5) *(Roughness comparisons) For every $z\le(N/(2L^C))^{\eta}$, the sums in (G4) restricted to the z-rough m, those all of whose prime factors are at least z, equal the unrestricted ones.*

For even N, $S\ge2C_2>1$, so $N^{1/2}L\le SNL^{-C}$ for large N, and (G0) gives $\sum_m a_m\Lambda(m)\ll SNL^{-C}$. This contradicts, for large N, every lower bound of the form $\sum_m a_m\Lambda(m)\ge\varepsilon\,\mathfrak{S}(N)N$ with fixed $\varepsilon>0$. Hence no argument that proves such a lower bound for every nonnegative sequence satisfying (G1)–(G5) relative to b' exists. Since $\Lambda(N-m)$ satisfies neither (G2) nor (G5) as stated (§9.3), the conclusion applies to arguments that use $\Lambda(N-m)$ only through (G1), (G3) and (G4), through bilinear estimates over factors in $(1,N^{\nu}]$, measured against b', which (G2) implies, and

through bounds for the sums of (G4) restricted to the z-rough m, $z \le (N/(2L^C))^\eta$, which (G4) and (G5) imply.

Construction. Fix α_* and ν_* with $\alpha < \alpha_* < \alpha'$ and $\nu < \nu_* < \eta$, and let $\varepsilon > 0$ and F_ε be given by Theorem 4.1 for the construction and the pair (ν_*, α_*); neither depends on N. Put $T = N/L^C$, $x_j = N/2^j$ for $0 \le j \le J = \lfloor C \log_2 L \rfloor$, and $B_j = (x_j/2, x_j] \cap [T, N)$. These sets partition $[T, N)$. The last one, $B_J = [T, x_J]$, is in general a proper subinterval of $(x_J/2, x_J]$, since $x_J/2 < T \le x_J$. Put

$$a_m = b'_m \big(1 + F_\varepsilon(v(m))\big) \quad (T \le m < N), \qquad a_m = b'_m \quad (1 \le m < T),$$

where $F_\varepsilon(v(m))$ is read as 0 unless m is squarefree and η-rough. Then $a \ge 0$, because $F_\varepsilon \ge -1$. For large N, $N^\alpha \le x_j^{\alpha_*}$ and $N^\nu \le x_j^{\nu_*}$ in every block, since $x_j \ge N/L^C$. Put $R = (N/(2L^C))^\eta$. For large N, $N^\nu < R$ and $R > w$, and the support of $a - b'$ consists of squarefree $m \ge T$ all of whose prime factors exceed R: if $m \in B_j$, its prime factors are at least $m^\eta > (x_j/2)^\eta \ge R$.

The only new analytic step is the passage from the baseline $b = 1$ of Theorem 4.1 to the presieved weight b'.

Lemma 9.3 (Transfer through the presieve)**.** *Let B_j be a block of the construction and $x = x_j$, and let $\mathcal{Q}$ be the set of squarefree $q \le x^{1/2}$ whose prime factors all exceed R (so $1 \in \mathcal{Q}$). For $q \in \mathcal{Q}$, an interval $I \subseteq B_j$ and $f \in \{1, \mu\}$ put*

$$\Sigma_f(q, I) = \sum_{\substack{m \in I \\ q \mid m}} f(m)\,(a_m - b'_m) = \sum_{\substack{m \in I \\ q \mid m}} f(m)\, b'_m\, F_\varepsilon(v(m)), \qquad \Sigma_f^\flat(q, I) = \sum_{\substack{m \in I \\ q \mid m}} f(m)\, F_\varepsilon(v(m)).$$

There is a number $c_N = \mathfrak{S}(N)\big(1 + O(e^{-L^{1/4}})\big)$, depending only on N, such that for all fixed $A, B > 0$ and uniformly in the block

$$\sum_{q \in \mathcal{Q}} \tau(q)^B \max_{I \subseteq B_j} \big|\Sigma_f(q, I) - c_N\, \Sigma_f^\flat(q, I)\big| \ll x L^{-A}.$$

Proof. Let $\mathcal{S}$ be the set of squarefree η-rough $m \in B_j$, and $C' = \sup |F_\varepsilon|$. Every $m \in \mathcal{S}$ has at most $1/\eta$ prime factors, all exceeding $R > w$, and for $q \in \mathcal{Q}$ at most x/q elements of $\mathcal{S}$ are multiples of q. We use $\sum_{q \le x} \tau(q)^B/q \ll L^{2^B}$, $K_w \ll \sqrt{L}$ and $\mathfrak{S}(N) \ll \log L$ throughout.

Step 1: the Euler factor and $(m, N) = 1$. On $\mathcal{S}$, $\prod_{p \mid m,\ p > w} \frac{p-1}{p-2} = 1 + O_\eta(R^{-1})$. At most $2/\eta$ primes $p > R$ divide N, and on $\mathcal{S}$ we have $\mathbf{1}_{(m,N)=1} = 1$ unless m is a multiple of one of them. So, uniformly in I, with an additive error,

$$\Sigma_f(q, I) = K_w\, \Sigma_f^+(q, I) + O\Big(\frac{K_w x}{qR} + K_w \sum_{p \mid N,\ p > R} \Big(\frac{x}{[q, p]} + 1\Big)\Big),$$

where $\Sigma_f^+(q, I) = \sum_{m \in I,\ q \mid m} f(m)\, F_\varepsilon(v(m))\, \mathbf{1}_{(N-m, P_w)=1}$. Weighted by $\tau(q)^B$ and summed over $\mathcal{Q}$, the error is $\ll x L^{O_B(1)} R^{-1} + x^{1/2} L^{O_B(1)} \ll x N^{-\eta/2}$.

Step 2: the presieve. With the weights of (5) for the primes dividing P_w, write $\theta^\pm = \theta^\pm_{P_w}$. Since $0 \le \theta^+ - \mathbf{1}_{(\cdot, P_w)=1} \le \theta^+ - \theta^-$ and $|F_\varepsilon| \le C'$,

$$\Sigma_f^+(q, I) = \sum_{e \mid P_w} \lambda_e^+\, \Sigma_f^\flat(q, I; e) + O\Big(C' \sum_{e \mid P_w} (\lambda_e^+ - \lambda_e^-)\, \#\{m \in \mathcal{S} : q \mid m,\ m \equiv N \bmod e\}\Big),$$

where $\Sigma_f^\flat(q, I; e)$ is $\Sigma_f^\flat(q, I)$ restricted to $m \equiv N \pmod e$. As $e \mid P_w$ is coprime to $2N$ and to every element of $\mathcal{S}$, this is a reduced residue class.

Step 3: the Bombieri–Vinogradov theorem in the largest prime. Let ψ be either $m \mapsto f(m)\,F_\varepsilon(v(m))$ or the indicator function of $\mathcal{S}$. We show that for every fixed A_1

$$\sum_{q\in\mathcal{Q}} \tau(q)^B \sum_{\substack{e\mid P_w\\ e<D}} \max_{I\subseteq B_j} \Big| \sum_{\substack{m\in I,\ q\mid m\\ m\equiv N \bmod e}} \psi(m) - \frac{1}{\varphi(e)} \sum_{\substack{m\in I\\ q\mid m}} \psi(m)\Big| \ll xL^{-A_1}. \tag{6}$$

Let $m \in \mathcal{S}$ with $q \mid m$. As $m > x/2 > q$, the cofactor m/q exceeds 1; write $m = qhP$ with P the largest prime factor of m/q. This factorisation is unique, $P > R$, and $m \equiv N \pmod e$ if and only if $P \equiv N\overline{qh} \pmod e$, again a reduced class. Fix q, h and a dyadic range $Y < P \le 2Y$. For real $t \in (Y, 2Y]$ let v_t be the increasingly ordered vector of the numbers $\log p/\log(qht)$, $p \mid qh$, and $\log t/\log(qht)$. Between the $O_\eta(1)$ values of t at which t passes a prime factor of q, the order of the coordinates is fixed and v_t moves monotonically along a line segment (its coordinates are affine in $1/\log(qht)$), by $O(1/\log x)$ in all. By Lemma 4.3, F_ε is in each dimension a finite sum of functions supported and Lipschitz on convex polytopes, and the conditions $qht \in I$, $t > P^+(h)$ (the largest prime factor of h) and the η-roughness of qht cut out intervals of t. The sign $f(qhP)$ does not depend on P. So $t \mapsto \psi(qht)$, with the squarefree condition dropped, has total variation $O(1)$ on $(Y, 2Y]$, uniformly in q, h and I, and partial summation bounds the difference between the sum of $\psi(qhP)$ over the primes $P \equiv c \pmod e$ in the range and $1/\varphi(e)$ times its sum over all primes in the range by

$$\ll \max_{y\le 2Y} \Big|\pi(y; e, c) - \frac{\pi(y)}{\varphi(e)}\Big|.$$

The $O(1)$ primes $P \mid q$, excluded by the squarefree condition, change this by $O(1)$ for each e, hence by $O(D)$ after summing over e. By the Bombieri–Vinogradov theorem [1], [19], [13, Chapter 17], in the form for $\pi(y; e, c)$, which follows from the form for $\psi(y; e, c)$ by partial summation, for every A_2 there is B_2 with

$$\sum_{e\le Y^{1/2}(\log Y)^{-B_2}} \max_{(c,e)=1} \max_{y\le 2Y} \Big|\pi(y; e, c) - \frac{\pi(y)}{\varphi(e)}\Big| \ll_{A_2} Y(\log Y)^{-A_2},$$

and $D \le Y^{1/2}(\log Y)^{-B_2}$ for $Y \ge R/2$ and large N. For fixed q and Y there are at most $x/(qY)$ values of h. Summing over h, over the $O(L)$ dyadic ranges and over $q \in \mathcal{Q}$ with the weight $\tau(q)^B$ gives (6) with $A_1 = A_2 - 1 - 2^B$, plus $O(D\,xL^{O(1)}/R)$ from the excluded primes, which is $\ll xN^{-\eta/2}$.

Step 4: assembly. The density $1/(p-1)$ at the primes $p \mid P_w$, all of which are at least 3, satisfies the dimension condition, so by (5) $\sum_e \lambda_e^{\pm}/\varphi(e) = \big(1 + O(e^{-L^{1/4}})\big)\prod_{p\mid P_w}\big(1 - \frac{1}{p-1}\big)$, and $K_w \prod_{p\mid P_w}\big(1 - \frac{1}{p-1}\big) = \mathfrak{S}(N)$ by the definition of K_w. Put $c_N = K_w \sum_e \lambda_e^+/\varphi(e)$. Inserting (6), for both choices of ψ, into Step 2 gives, uniformly in I,

$$K_w\,\Sigma_f^+(q, I) = c_N\,\Sigma_f^\flat(q, I) + O\Big(\mathfrak{S}(N)\,e^{-L^{1/4}}\,\frac{x}{q}\Big) + K_w\,\mathcal{E}_q, \qquad \sum_{q\in\mathcal{Q}} \tau(q)^B\,\mathcal{E}_q \ll xL^{-A_1}.$$

With Step 1 and $A_1 = A + 1$, this proves the lemma. □

Proof of Proposition 9.2. We use the notation of Lemma 9.3, with $x = x_j$ in the block B_j. Lemma 4.4 and its proof hold with $(x/2, x]$ replaced by any subinterval $(cx, c'x]$, $1/2 \le c \le c' \le 1$, because [5, Lemma 5.11], on which the proof rests, is uniform in c and c'; the integral of Lemma 4.4 (in its variable w and with its $\lambda = \lambda_1 + \cdots + \lambda_t$; these are not the presieve level and the sieve weights of this section) then runs over $[1 - \lambda + \log c/\log x,\ 1 - \lambda + \log c'/\log x]$, and the integrand is still zero. The same holds for the case $d = 1$ treated in Step 4 of §4, with $\mathrm{Li}(c'x) - \mathrm{Li}(cx)$ in place of $\mathrm{Li}(x) - \mathrm{Li}(x/2)$. Every block, the truncated B_J included, has the form $(cx_j, c'x_j] \cap \mathbb{Z}$ with $1/2 \le c \le c' \le 1$, and so has every interval contained in a block.

- *(G0).* A prime p is squarefree and η-rough with $v(p) = (1)$, and $F_\varepsilon(1) = -1$, so $a_p = 0$ for primes $p \geq T$. On proper prime powers $F_\varepsilon(v(m))$ is read as 0, so there $a_m = b'_m \leq K_w(1 + O(1/w)) \ll \sqrt{L}$; as $\sum_{p^k<N,\ k\geq 2} \log p \ll N^{1/2}$, they contribute $\ll N^{1/2}L$. For the primes $w < p < T$, $a_p = b'_p \leq K_w(1 + O(1/w))\mathbf{1}_{(N-p,P_w)=1}$. With the weights λ_e^+ and the Bombieri–Vinogradov theorem at level $D \leq T^{1/2}L^{-B_2}$,

$$\begin{aligned}\#\{w < p < T : (N-p, P_w) = 1\} &\leq \sum_{e|P_w} \lambda_e^+ \,\pi(T; e, N) + O(wD) = \pi(T) \sum_{e|P_w} \frac{\lambda_e^+}{\varphi(e)} + O(TL^{-2}) \\ &= \frac{\mathfrak{S}(N)}{K_w}\,\pi(T)\big(1 + o(1)\big) + O(TL^{-2}).\end{aligned}$$

  The primes $p \leq w$ contribute $O(K_w w \log w)$. Hence $\sum_{p<T} a_p \log p \ll \mathfrak{S}(N)\, NL^{-C}$, and (G0) follows.
- *(G1).* Let $d \leq N^{1/2}L^{-C/2}$ and let $I \subseteq [1, N)$ be an interval. The terms with $m < T$ vanish, so $\sum_{m\in I,\ d|m}(a_m - b'_m) = \sum_j \Sigma_1(d, I \cap B_j)$. This is 0 unless d is squarefree with all prime factors greater than R, and then $d \in \mathcal{Q}$ for every j, since $x_j \geq N/L^C$ gives $d \leq x_j^{1/2}$. By Lemma 9.3 with $f = 1$, and since $c_N \ll \log L$, it suffices to show $\sum_{d\leq x_j^{1/2}} \tau(d)^B \max_I |\Sigma_1^\flat(d, I \cap B_j)| \ll x_j L^{-A-1}$. Here $\Sigma_1^\flat(d, \cdot)$ is the Type I sum of the sequence $w_n = F_\varepsilon(v(n))$ of Theorem 4.1 at scale x_j, over the multiples of d in an interval of $(x_j/2, x_j]$. Condition (I) of Theorem 4.1, with $\max(A + 2, B)$ in place of its B, gives the bound, since $\log x_j \geq L/2$. Summing over the $J + 1 \ll \log L$ blocks gives (G1).
- *(G2).* Every element of the support of $a - b'$ has all its prime factors greater than $R > N^\nu$. So no $\ell \in (1, N^\nu]$ divides an element of the support, and every term of the sum in (G2) vanishes.
- *(G3).* Let $q \leq N^\alpha$. As in (G1), $\sum_{m<N,\ q|m} \mu(m)(a_m - b'_m) = \sum_j \Sigma_\mu(q, B_j)$, and this is 0 unless $q \in \mathcal{Q}$: a non-squarefree q gives 0 through μ, and a q with a prime factor at most R divides no element of the support. Let $q = p_1 \cdots p_t \in \mathcal{Q}$. Then $t \leq m^* = \max\{j : j\eta < \alpha'\}$: otherwise $t\eta \geq (m^* + 1)\eta \geq \alpha'$, and $q > R^t \geq (N/(2L^C))^{\alpha'} > N^\alpha$ for large N, because $\alpha < \alpha'$ strictly. (At $(1/18, 1/6)$, $(1/20, 1/5)$ and $(1/30, 4/15)$, $(m^* + 1)\eta = \alpha'$ exactly, so these cases use the strictness.) Also $q \leq N^\alpha \leq x_j^{\alpha_*} \leq x_j^{1/2}$ in every block. By Lemma 9.3 with $f = \mu$ and $I = B_j$ (the weight $\tau(q)^B \geq 1$ may be dropped), and since $c_N \ll \log L$, it suffices to show $\sum_q |\Sigma_\mu^\flat(q, B_j)| \ll x_j e^{-(\log x_j)^{1/5}}$, the sum over $q \in \mathcal{Q}$ with $q \leq N^\alpha$. In the block of scale $x = x_j$, every element of the support has all prime factors at least $(x/2)^\eta$, so $\Sigma_\mu^\flat(q, B_j)$ is 0 unless every $p_i \geq (x/2)^\eta$. In that case, for $t \geq 1$, $\Sigma_\mu^\flat(q, B_j)$ is the sum S_q^μ of Lemma 4.4 at scale x, with α_* in place of α, restricted to B_j, and Lemma 4.4 in the subinterval form above gives $\ll (x/q)\, e^{-(\log x)^{1/5}}$, because $T_t(F_\varepsilon)$ vanishes at every t-tuple with total at most α'. For $q = 1$ the case $d = 1$ of Step 4 of §4, in the subinterval form, gives $(\mathrm{Li}(c'x) - \mathrm{Li}(cx))\, T_0(F_\varepsilon) + O(xe^{-(\log x)^{1/5}})$, and $T_0(F_\varepsilon) = 0$. Summing over q costs at most a factor $\sum_{t\leq m^*} \big(\sum_{R<p\leq N} 1/p\big)^t \ll 1$, since $\sum_{R<p\leq N} 1/p = \log(1/\eta) + o(1)$. Summing over the blocks gives (G3).
- *(G4).* Let $I \subseteq [1, N)$ be an interval. As in (G1), $\sum_{m\in I}(a_m - b'_m)\mu(m) = \sum_j \Sigma_\mu(1, I \cap B_j)$. By Lemma 9.3 with $q = 1$ and $f = \mu$ it suffices that $\Sigma_\mu^\flat(1, (cx_j, c'x_j]) \ll x_j e^{-(\log x_j)^{1/5}}$ uniformly for $1/2 \leq c \leq c' \leq 1$. This is the subinterval form of the case $d = 1$: the sum is $(\mathrm{Li}(c'x_j) - \mathrm{Li}(cx_j))\, T_0(F_\varepsilon) + O(x_j e^{-(\log x_j)^{1/5}})$, and $T_0(F_\varepsilon) = 0$. So (G4) holds for $f = \mu$; and $a - b'$ is supported on squarefree integers, where $\lambda = \mu$.
- *(G5).* The support of $a - b'$ consists of integers whose prime factors all exceed R. They are z-rough for every $z \leq R$, so restricting the sums in (G4) to z-rough m changes nothing. □

Construction (η, α')	Twisted level	Type II range	Separate implementations
$(1/11, 1/10)$	$\alpha < 1/10$	$\nu < 1/11$	1
$(1/14, 1/8)$	$\alpha < 1/8$	$\nu < 1/14$	1
$(1/18, 1/6)$	$\alpha < 1/6$	$\nu < 1/18$	2
$(1/20, 1/5)$	$\alpha < 1/5$	$\nu < 1/20$	2
$(1/30, 1/4)$	$\alpha < 1/4$	$\nu < 1/30$	1
$(1/30, 4/15)$	$\alpha < 4/15$	$\nu < 1/30$	1

TABLE 4. The cases of Proposition 9.2. Each uses the certified function of the stated construction and its lift (Theorem 4.1). The last column is the number of separately written implementations that certify the construction (§5.6).

The base case $(1/11, 1/11)$ of Remark 6.1 needs no row: its ranges $\alpha < 1/11$ and $\nu < 1/11$ lie inside those of $(1/11, 1/10)$, so it follows as a corollary of that case, without a separate certificate.

The reflection argument for each case was audited independently, in an earlier written form: the audit reconstructed the passage from the certified function and its lift to (G0)–(G5), confirmed it, and required corrections, which the proof above incorporates; for $(1/20, 1/5)$ it also covered the cap computation below. These audits check the reflection, not the constructions. The constructions' certification is that of §5.6: $(1/18, 1/6)$ and $(1/20, 1/5)$ are certified by two separately written implementations each and the others by one; for $(1/30, 1/4)$ and $(1/30, 4/15)$ that one is the second-order engine, whose code was audited, and no second certificate exists.

9.3. **What the proposition does and does not say.**

- *Strict ranges only.* Each case holds for every fixed α below its level and every fixed ν below its η. No endpoint is claimed. No threshold beyond the last case is asserted, and none follows from Conjecture 8.2, which concerns a family of constructions, not Goldbach's problem.
- *Type II is stronger than the truth's.* (G2) is an exact identity for arbitrary coefficients, and the true sequence fails that form. Take a prime $\ell \le 2L$ with $\ell \nmid N$. Such a prime exists for large N: otherwise $\prod_{p \le 2L} p$ divides N, whereas its logarithm is $2L + o(L) > L$ by the prime number theorem. Set $\xi = \mathbf{1}_{\{\ell\}}$, and $\kappa_k = \operatorname{sgn}(\Lambda(N - \ell k) - b'_{\ell k})$ for $N/2 < \ell k < N$, $\kappa_k = 0$ otherwise. The sum is then $\sum_{N/2 < \ell k < N} |\Lambda(N - \ell k) - b'_{\ell k}| \asymp N/\ell$, by the prime number theorem in progressions and Lemma 9.1, so the true sequence fails even Ford and Maynard's bounded-coefficient form (II) with $\theta = 0$ on the block $(N/2, N)$, for every $\nu > 0$. So (G2) is not an input that the truth provides. What it shows is that every weaker small-factor bilinear bound measured against b' also holds for the reflected sequence, whatever coefficient class the true sequence satisfies, because the reflected sums vanish for all coefficients.
- *Type I is shared only up to a common level.* (G1) reaches $d \le N^{1/2} L^{-C/2}$. For $\Lambda(N - m)$, the sums over the multiples of d are controlled by the Bombieri–Vinogradov theorem (the weight $\tau(d)^B$ is removed by the Cauchy–Schwarz inequality; for $(d, N) > 1$ only powers of primes dividing (d, N) occur, as after Lemma 9.1) only for $d \le N^{1/2} L^{-B'}$, with some $B' = B'(A, B)$. That b' satisfies the matching estimate, $\sum_{d \le N^{1/2}} \tau(d)^B \max_I \big|\sum_{m \in I,\ d \mid m} b'_m - \mathbf{1}_{(d,N)=1} |I|/\varphi(d)\big| \ll NL^{-A}$, follows from the fundamental lemma as in the proof of Lemma 9.1, with the Euler factor of b'_m expanded as a divisor sum; we do not write this out. The two sequences are therefore compared on $d \le \min(N^{1/2} L^{-C/2}, N^{1/2} L^{-B'})$; when $C < 2B'$ the longer range of the reflected sequence is an extra property, like (G2) and (G3).
- *Hooley's condition and the parity sums are granted.* (G3) is not known for $\Lambda(N - m)$ at any fixed power level, and (G4) is not known for it at the precision NL^{-A}, even on $[1, N)$ with $f = \mu$. There the b'-part $\sum_{m < N} b'_m \mu(m)$ is $\ll NL^{-A}$ (by the presieve weights of

(5), with the Euler factor of b'_m expanded as a divisor sum, and the Bombieri–Vinogradov theorem for μ to the moduli below D; we omit the details), so (G4) amounts to a bound for the Möbius function on the shifted primes $N-p$. What is known for the true sequence is weaker and recent. Lü, Wang and Wu give $\sum_{p\le x}\mu(p+2)=o(\pi(x))$ [14]. Guo deduces from their arguments $\sum_{p<N}\mu(N-p)=o(\pi(N))$, and the same for λ, with the rate $O(\pi(N)/\sqrt{\log\log N})$ [9]. He assumes their arguments as written and, for general N, a zero-free region for Dirichlet L-functions; when $\mathrm{rad}(N)\le(\log N)^{O(1)}$ he does not need the second assumption. Both works are recent preprints. (G5) is, like (G2), an exact identity for the reflected sequence, and the true sequence fails it as written: for $f=\lambda$, let 2^k be the largest power of two strictly below N, set $m=N-2^k$, and take $I=\{m\}$ and $z=3\le R$. Then $0<m\le N/2$ is even, $b'_m=0$, and $\Lambda(N-m)=\log 2>0$. The unrestricted singleton sum is nonzero, while its restriction to z-rough integers is zero. Including (G3)–(G5) only strengthens the conclusion.

- *Huang and Li's criterion is untouched.* Their twisted hypothesis concerns the true Goldbach sequence, subtracts the full initial-interval sum divided by $\varphi(q)$, and takes maxima over initial intervals and reduced residue classes. Its level is $N^{1-\theta}$, paired with prime Elliott–Halberstam at $N^{\theta}(\log N)^{2A+8}$, in their notation (§1). In the specialization where the latter is supplied by Bombieri–Vinogradov, $\theta<1/2$, so the twisted level exceeds $1/2$. Proposition 9.2 grants (G3) only below $4/15$, measured against b', and establishes no version of their criterion.
- *Plain and switching versions.* The proposition asserts no explicit cap $a_m\le Kb'_m$; the bound $|F_\varepsilon|\le C'$ of §4, Step 3, gives one only with an inexplicit K. Explicit caps matter for Chen-type switching arguments. They need an upper enclosure of f, and the certificates of §5 bound f on non-prime vectors only from below. For $(1/20,1/5)$, an upper-bound pass added to the code of the certificate of record and run at its settings ($M=7200$; its lower-bound outputs agree with the stored certificate field by field) gives, in exact rational arithmetic, $f(u)\le 2011282180584704051/(5\cdot 10^{17})<4.022565$ for every non-prime η-rough vector $u\ne E$. With $|f(1)|\ge 527670700557531721/(5\cdot 10^{17})$, this gives $F\le 2011282180584704051/527670700557531721<3.811624$. Since $f_\varepsilon=f+O(\varepsilon)$ off the mask and $f_\varepsilon=0$ on it (§4, Step 3), $F_\varepsilon<4$ for ε small enough, and then the reflected sequence satisfies $a_m\le 5b'_m$. The same computation gives $F>3.8$ on an open region of two-component vectors. Take an open box $\mathcal{V}$ inside it whose closure contains no vector with a component $1/2$; for small ε, $\mathcal{V}$ is unmasked and $F_\varepsilon>3$ on it. The reflected indices meet $\mathcal{V}$ with positive b'-mass for every large N: Steps 1–4 of the proof of Lemma 9.3, with $q=1$, $f=1$ and F_ε replaced by the indicator function of $\mathcal{V}$ (a convex polytope, so Step 3 applies unchanged), give $\sum b'_m=c_N\,\#\{m\}+O(NL^{-2})$ over the squarefree $m\in B_0$ with $v(m)\in\mathcal{V}$, and there are $\gg N/L$ such m. So $a_m\le 4b'_m$ fails for this function. The pass is one implementation; an independent audit read its code and reran it at $M=3600$ and 7200 (§10). No cap is claimed for the other cases.

10. Computations

The theorems cite stored outputs. Each margin in Table 2 comes from a certificate of record, run by the author at the resolution given there (for $(1/18,1/6)$ and $(1/20,1/5)$ there are two, from separate implementations; the full runs of both separate implementations were repeated with corrected programs, each reproducing every printed line apart from timings, §5.5). Its output is part of the supplement, with the logs that the supplement's `README.md` lists. Each JSON output records the enclosures of $T_0(c)$, $T_0(g)$, t^* and $f(1)$, the two positivity bounds, the margin as an exact fraction, the verdict and the SHA-256 of the input function. The first-order JSON outputs also record the range of each diagonal and each forced function. The second-order outputs record

the SHA-256 of the engine and of the arithmetic module; their diagonal bounds are printed in the logs, rounded to the nearest fourth decimal, and the bound 0.825079 for $(1/20, 1/5)$ is taken from the audit's replay (`audit/p5/p5_diagnostics_M7200.json`). The separate implementations write text outputs, with enclosures rounded outward to five or six decimals and the margin as a binary64 number; they read their constants from the TSV files beside them and record no hash. The supplement is the `anc/` directory of the paper's record; its `README.md` describes every directory, and `SHA256SUMS.txt` lists every file. Single-threaded, the second-order runs take about 15 seconds for $(1/18, 1/6)$, five minutes for $(1/20, 1/5)$, 4.6 hours for $(1/30, 1/4)$ and nearly 18 hours for $(1/30, 4/15)$; Table 5 lists the full runs of record, together with the upper-bound run of §9. A reader can check a certificate by rerunning it in full, or more cheaply by checking the code against the text and running the small-resolution consistency runs below. Of the other certificate runs whose margins §§5 and 6 quote, only the small-resolution consistency runs below have stored outputs in the supplement; the others (the margins of $(1/11, 1/10)$ at $M = 26{,}400$ and 52,800, those of the exact first-order certificate of $(1/18, 1/6)$ at $M = 108{,}000$ and 144,000, and the variant of $(1/14, 1/8)$ with 12 constants) are reported for context, as is the binary64 computation with FFT convolutions. The audit's replay of $(1/30, 1/4)$ at $M = 1920$, whose margin is negative (§5.5), is not in the supplement either.

What the supplement provides for each certificate of record (Table 5) is:

- the code, the input function and the stored outputs, with SHA-256 hashes;
- the exact command and environment of the full run, its elapsed time (rounded) and peak memory where it was recorded;
- a consistency run at small resolution, single-threaded, with the settings of the full run: the exact first-order certificates at $M = 13{,}200$, 53,760 and 7200 for $(1/11, 1/10)$, $(1/14, 1/8)$ and $(1/18, 1/6)$; the second-order engine at $M = 3600$ for $(1/20, 1/5)$, together with its upper-bound pass, and at $M = 960$ for $(1/30, 1/4)$ and $(1/30, 4/15)$; and the separate implementations at $M = 144{,}000$ for $(1/18, 1/6)$ and at $M = 36{,}000$, with the coarse resolution equal to M, for $(1/20, 1/5)$. A supplied script compares its outputs field by field with stored small-resolution outputs. The second-order run of $(1/18, 1/6)$ takes about 15 seconds and has no small-resolution run; `run_all.sh.txt` and `compare.py` do not cover it. It is checked by rerunning in full the command that the supplement's `README.md` gives and comparing the output with the stored one, which it reproduces in every value field.

The consistency runs take about an hour and a quarter in all on one core. They check that the code and inputs as shipped reproduce the stored computation. They are not what the theorems cite: at these resolutions several margins are negative, and the theorems rest on the full runs of Table 5. They are the author's own reproduction of the computations against stored outputs, not an independent audit; in particular, the run for $(1/30, 4/15)$ at $M = 960$ runs that certificate end to end with the settings of its full run, which no independent audit did (§5.5). Every code requires η, α' and every bin edge to lie on the grid $\{n/M\}$, and the second-order engine requires M even. The exact $(1/18, 1/6)$ certificate also requires $36 \mid M$ and both block counts to divide $M/9$ ($900 \mid M$ with 50 blocks); the separate $(1/18, 1/6)$ implementation requires $36 \mid M$ and its block parameter Q to divide $M/18$ ($72 \mid M$ at the value $Q = 4$ used); the separate $(1/20, 1/5)$ implementation requires $60 \mid M$ and $60 \mid M_c \mid M$.

The second-order runs. Each takes the resolution M, the number NB of threshold blocks and the function, with the profile, the parity direction and the block parameters set through the environment: `KPROF` and `KW` set the profile and its weights, `GBIN` or `GVEC` the parity direction g, `TOADJ` and `F1ADJ` switch on the adjoint enclosures of T_0 and $f(1)$, `NBQ` and `NBQLO` are the block counts of §5, `QOUTER` selects block-by-block evaluation, and `QADAPT` gives, for one extended class, the target, the splitting factor and the depth of the adaptive partition. The runs of record are:

Case	Implementation	M	elapsed time	peak memory
$(1/11, 1/10)$	exact integer, first order	79,200	about 20 s†	514 MiB
$(1/14, 1/8)$	exact integer, first order	107,520	about 2 min	1085 MiB
$(1/18, 1/6)$	exact integer, first order	216,000	about 14 min	2199 MiB
$(1/18, 1/6)$	validated binary64, first order (separate)	288,000	about 11 min	813 MiB
$(1/18, 1/6)$	second order	1800	about 15 s	not recorded
$(1/20, 1/5)$	second order	7200	about 5 min	not recorded
$(1/20, 1/5)$	validated binary64, first order (separate)	432,000*	about 2.4 h‡	3606 MiB
$(1/20, 1/5)$	second order, upper-bound pass (§9)	7200	about 4 min	not recorded
$(1/30, 1/4)$	second order	7680	about 4.6 h	509 MiB
$(1/30, 4/15)$	second order	23,040	about 17.8 h	889 MiB

TABLE 5. The full runs of record, single-threaded, on the machine that produced the stored outputs. Elapsed times are wall-clock times as recorded by the run (by GNU `time` where available, otherwise by the script's own timer), rounded: that machine's wall clock is unreliable over short runs: the long runs' wall-clock and CPU times, where both were recorded, agree to about ten per cent (the run marked ‡ apart), but in the run of $(1/11, 1/10)$ the script's time stamps are not monotone and the recorded 21 s of wall-clock time is less than its 34 s of CPU time. Peak memory is the maximum resident set size, where it was recorded. *Some extended classes on the coarser grid $\{n/72{,}000\}$ (§5.2). †As recorded; the run used 34 s of CPU time. ‡CPU time; the run's own elapsed time, 3.2 h, includes a period in which the machine paused the job.

- $(1/18, 1/6)$: $M = 1800$, $NB = 10$, engine `cert2_v2.py` (`045cc3ae`), with its defaults (all profiles, refined blocks, $g = e_{(0,3,1)}$); the lift uses the exact first-order margin of this case, not this run.
- $(1/20, 1/5)$: `KPROF=3`, $M = 7200$, $NB = 40$, engine `cert2_v2.py` (`045cc3ae`); the parity direction is $g = e_{(0,3,4)}$, the engine's default for this case, in which `GBIN` is the number of bins, 4.
- $(1/30, 1/4)$: `KW=2:1,3:-1 GVEC=2:1,3:1 TOADJ=1 NBQ=32 NBQLO=64`, $M = 7680$, $NB = 80$, development engine `5c779c07`.
- $(1/30, 4/15)$: engine `cert2_f2.py` (`0ba72980`), $M = 23{,}040$, $NB = 80$, with `KW=2:1,3:-1 GVEC=2:1,3:1 TOADJ=1 F1ADJ=1 NBQ=32 NBQLO=64 QOUTER=1 QADAPT=2,1:-1.0147:8:2`.

All four use the same exact arithmetic module (`dpl.py`, `b7f331da`). The inputs of the $(1/30, 4/15)$ run were hashed and frozen before it was launched.

The first-order runs. Each takes the resolution M and the input function; the certificates of record for $(1/11, 1/10)$ and $(1/14, 1/8)$ also take the list of nonzero classes and the parity direction on the command line, and that of $(1/18, 1/6)$ the numbers of threshold and σ_β blocks (50 and 50). Their exact commands are in the supplement. The separate implementations use the arithmetic of §5.2. The separate $(1/20, 1/5)$ run is `p5rig.py 432000 cert 1 72000 16`: the fine resolution is $M = 432{,}000$, and the extended strata with $j_\beta \in \{1, 2\}$ and $k_\beta \in \{1, 2\}$, which carry no parity term, are computed at the coarser resolution $M_c = 72{,}000$, with σ_β split into 16 groups. The program's default is $M_c = M$, so M_c must be given.

The upper-bound run of §9. For $(1/20, 1/5)$, `cert2_caps.py` (`081b0528`) is `cert2_v2.py` with an upper-bound pass that mirrors every lower-bound step in the same arithmetic. Run on the function of record at the settings of record (`KPROF=3`, $M = 7200$, $NB = 40$), it reproduces every lower-bound field of the certificate of record. It adds $\sup f \le 2011282180584704051/(5 \cdot 10^{17}) < 4.022565$, enclosed in the extended class $(0, 1)$, and the free bound $f < 3.037643$. With $|f(1)| \ge$

$527670700557531721/(5 \cdot 10^{17})$ these give the ratio quoted in §9. Its audit read the added code against the text and replayed the run at $M = 7200$ (every output field equal) and at $M = 3600$. Theorem B. Its bounds come from exact-integer cell-convolution computations at $\nu = 3/20$, $1/7$ and $1/8$ (mesh 1/16,000, 1/9800 and 1/8000), made during the audit of that theorem with separately written scripts (§5.5), and at $\nu = 0.116$ (mesh 1/32,000), made after the audit by `independent_p1_k8.py`, a copy of the audit's program with the term in I_8 added; with that term set to zero the copy reproduces the stored exact outputs at 3/20, 1/7 and 1/8 byte for byte. The scripts and their outputs, together with the exact computation behind Remark 6.1, are in the supplement's `thmB/`, and the supplement's `README.md` gives the commands that produced them. The run at $\nu = 3/20$ with mesh 1/8000, also in `thmB/`, is by the same program and gives a weaker bound; no exact computation of the bounds at $\nu = 3/20$, $1/7$ and $1/8$ by a second, separately written program is reported here. At $\nu = 0.116$ a second, separately written program, `k8_bound.py`, also written and run after the original audit, applies the same scheme with mesh 1/1,024,000 and gives a lower bound of at least 0.0354. A floating-point integral recurrence agrees at $\nu = 3/20$, $1/7$ and 0.116 (§7).

Appendix A. The class constants and the certified bounds

Each case uses the class constants c below, at their exact binary64 values, shown to six decimals and given exactly as C99 hexadecimal floating-point literals, with $c = 0$ on every other class. These are the input constants. The certified function uses $c^* = c + t^* g$, where g is the parity direction of Table 2 and t^* is the exact solution of $T_0(c + tg) = 0$; the certificate encloses t^* in the interval given in §6 and bounds f on the whole segment of constants $c + tg$ over that interval. Every coefficient of g is 1, so the constant on each class of g is $c + t^*$. Below each table are the exact bounds of the certificate of record: U is the upper bound for $f(1)$, L_{free} and L_{ext} are the lower bounds on free and on extended non-prime vectors, and the margin is $1 - \min(L_{\text{free}}, L_{\text{ext}})/U$.

P3a, case $(1/11, 1/10)$**: 11 constants.**

class	value	class	value	class	value
$(0,3,3)$	−1.000000	$(0,4,2)$	−1.000000	$(0,4,3)$	−1.000000
$(0,5,1)$	−1.000000	$(0,5,2)$	−0.512163	$(0,5,3)$	0.850519
$(0,6,1)$	0.092611	$(0,6,2)$	−0.185014	$(0,6,3)$	0.016115
$(1,2,3)$	−0.490802	$(2,3,3)$	−1.000000		

Exact values, in the order of the table: `-0x1.0000000000000p+0`, `-0x1.0000000000000p+0`, `-0x1.0000000000000p+0`, `-0x1.0000000000000p+0`, `-0x1.063a49c5bb725p-1`, `0x1.b3773feddc956p-1`, `0x1.7b55643d2ae8dp-4`, `-0x1.7ae8c1012ff98p-3`, `0x1.08066072f729bp-6`, `-0x1.f694cdd1717c1p-2`, `-0x1.0000000000000p+0`.

Certificate of record at $M = 79{,}200$:

U	$-547562092911977041/500000000000000000$
L_{free}	$-501134667639647357/500000000000000000$
L_{ext}	$-1006101272505629611/1000000000000000000$
margin	$89022913318324471/1095124185823954082$

P3b, case $(1/14, 1/8)$**: 19 constants.**

class	value	class	value	class	value
$(0,3,1)$	−0.877983	$(0,4,1)$	−1.000000	$(1,2,1)$	−0.608128
$(1,3,1)$	−0.608128	$(1,4,1)$	0.185643	$(2,2,1)$	0.278103
$(2,3,1)$	−0.993094	$(2,4,1)$	0.745909	$(3,2,1)$	−1.000000
$(3,3,1)$	0.518544	$(3,4,1)$	−0.509728	$(4,2,1)$	0.723092
$(5,5,1)$	4.000000	$(8,3,1)$	4.000000	$(10,2,1)$	4.000000
$(11,1,1)$	−1.000000	$(12,0,0)$	−1.000000	$(12,1,1)$	−1.000000
$(13,0,0)$	−1.000000				

Exact values, in the order of the table: `-0x1.c1870110dd9d1p-1`, `-0x1.0000000000000p+0`, `-0x1.375c8443d49bap-1`, `-0x1.375c8443d49bbp-1`, `0x1.7c32556af388cp-3`, `0x1.1cc6fab346f22p-2`, `-0x1.fc76d155926d7p-1`, `0x1.7de7d416055e5p-1`, `-0x1.0000000000000p+0`, `0x1.097e9100b3268p-1`, `-0x1.04fb0f9f6ab84p-1`, `0x1.7239188c58910p-1`, `0x1.0000000000000p+2`, `0x1.0000000000000p+2`, `0x1.0000000000000p+2`, `-0x1.0000000000000p+0`, `-0x1.0000000000000p+0`, `-0x1.0000000000000p+0`, `-0x1.0000000000000p+0`.

Certificate of record at $M = 107{,}520$:

U	$-538687099033821133/500000000000000000$
L_{free}	$-1004526547095420909/1000000000000000000$
L_{ext}	$-507637385794574409/500000000000000000$
margin	$2822701203567884/48971554457620103$

P4, case $(1/18, 1/6)$**: 13 constants.**

class	value	class	value	class	value
$(0,3,1)$	-0.926504	$(0,4,1)$	-1.000000	$(2,2,1)$	-0.912216
$(2,3,1)$	-0.373099	$(2,4,1)$	-0.912216	$(3,2,1)$	-0.929777
$(3,3,1)$	-0.543236	$(4,1,1)$	-1.000000	$(4,2,1)$	-0.654841
$(5,1,1)$	-0.786825	$(5,2,1)$	0.543837	$(6,1,1)$	0.519211
$(6,2,1)$	-0.225199				

Exact values, in the order of the table: `-0x1.da5eae936606bp-1`, `-0x1.0000000000000p+0`, `-0x1.d30e035d98884p-1`, `-0x1.7e0da01ee8c8bp-2`, `-0x1.d30e035d98898p-1`, `-0x1.dc0bb0aad550ep-1`, `-0x1.1622ffda543eep-1`, `-0x1.0000000000000p+0`, `-0x1.4f4757419d8f2p-1`, `-0x1.92daaa78634dfp-1`, `0x1.1671dac592bb2p-1`, `0x1.09d6001989713p-1`, `-0x1.cd355ef3df17fp-3`.

Certificate of record at $M = 216{,}000$:

U	$-524907394385979697/500000000000000000$
L_{free}	$-100353758199060953/100000000000000000$
L_{ext}	$-1019266358936212957/1000000000000000000$
margin	$30548429835746437/1049814788771959394$

P5, case $(1/20, 1/5)$**: 32 constants.**

class	value	class	value	class	value
$(0,3,3)$	-1.000000	$(0,3,4)$	-0.163554	$(0,4,2)$	-1.000000
$(1,2,3)$	-1.000000	$(1,2,4)$	-0.849365	$(2,2,3)$	-1.000000
$(2,2,4)$	-1.000000	$(2,3,1)$	-1.000000	$(2,3,2)$	-1.000000
$(3,2,2)$	-1.000000	$(3,2,3)$	-0.145241	$(3,2,4)$	-0.547640
$(3,3,2)$	-1.000000	$(3,3,3)$	-1.000000	$(4,1,3)$	-1.000000
$(4,1,4)$	-1.000000	$(4,2,2)$	-0.155968	$(4,2,3)$	-1.000000
$(4,2,4)$	0.329228	$(4,3,1)$	-1.000000	$(5,1,2)$	-1.000000
$(5,1,3)$	0.387196	$(5,1,4)$	-1.000000	$(5,2,2)$	-1.000000
$(5,2,3)$	-0.540169	$(5,2,4)$	0.000336	$(6,1,3)$	2.657027
$(6,1,4)$	-0.388139	$(6,2,1)$	-0.150240	$(6,2,2)$	1.396150
$(6,2,3)$	-0.936310	$(7,1,2)$	-1.000000		

Exact values, in the order of the table: `-0x1.0000000000000p+0`, `-0x1.4ef5850c272f8p-3`, `-0x1.0000000000000p+0`, `-0x1.0000000000000p+0`, `-0x1.b2e001e3f73f0p-1`, `-0x1.0000000000000p+0`, `-0x1.0000000000000p+0`, `-0x1.0000000000000p+0`, `-0x1.0000000000000p+0`, `-0x1.0000000000000p+0`, `-0x1.2974459faaba9p-3`, `-0x1.186440d333e64p-1`, `-0x1.0000000000000p+0`, `-0x1.0000000000000p+0`, `-0x1.0000000000000p+0`, `-0x1.0000000000000p+0`, `-0x1.3f6c3524b727ap-3`, `-0x1.0000000000000p+0`, `0x1.51213cd470a6bp-2`, `-0x1.0000000000000p+0`, `-0x1.0000000000000p+0`, `0x1.8c7d0241f459ep-2`, `-0x1.0000000000000p+0`, `-0x1.0000000000000p+0`, `-0x1.1490fea960ef7p-1`, `0x1.5fe2a04cdb06bp-12`, `0x1.5419799beb1b8p+1`, `-0x1.8d74343d7c585p-2`, `-0x1.33b0fc77e9eb1p-3`, `0x1.656a14697c327p+0`, `-0x1.df640e9553692p-1`, `-0x1.0000000000000p+0`.

Certificate of record at $M = 7{,}200$:

U	$-527670700557531721/500000000000000000$
L_{free}	$-1014556515501463499/1000000000000000000$
L_{ext}	$-252862019809166663/250000000000000000$
margin	$40784885613599943/1055341401115063442$

P6, case $(1/30, 1/4)$**: 45 constants.**

class	value	class	value	class	value
$(0,3,2)$	-0.309434	$(0,3,3)$	-0.930000	$(1,2,2)$	-0.134473
$(1,2,3)$	-1.871514	$(2,2,1)$	-0.995000	$(2,2,2)$	-0.995000
$(2,2,3)$	-0.995000	$(2,3,1)$	-0.780040	$(2,3,2)$	-0.780040
$(3,1,1)$	-0.995000	$(3,1,2)$	-0.995000	$(3,1,3)$	-0.995000
$(3,2,1)$	-0.728878	$(3,2,2)$	-0.311136	$(3,2,3)$	-0.758914
$(3,3,1)$	-0.475173	$(4,1,1)$	-0.995000	$(4,1,2)$	-0.995000
$(4,1,3)$	-0.995000	$(4,2,1)$	-0.808188	$(4,2,2)$	-0.808188
$(4,2,3)$	-0.808188	$(5,0,0)$	-0.995000	$(5,1,1)$	-0.995000
$(5,1,2)$	-0.995000	$(5,1,3)$	-0.340671	$(5,2,1)$	-0.563416
$(5,2,2)$	-0.154190	$(5,2,3)$	-0.101685	$(6,0,0)$	1.810431
$(6,1,2)$	-0.562839	$(6,2,2)$	0.469814	$(6,2,3)$	0.361605
$(7,0,0)$	-0.995000	$(7,1,1)$	-0.995000	$(7,1,2)$	-0.995000
$(7,1,3)$	0.065048	$(7,2,3)$	0.481914	$(8,0,0)$	-0.828270
$(8,1,1)$	-0.022034	$(8,1,2)$	-0.995000	$(10,0,0)$	1.289040
$(10,1,3)$	0.058522	$(11,0,0)$	1.078571	$(12,0,0)$	0.301125

Exact values, in the order of the table: `-0x1.3cdc443914f48p-2`, `-0x1.dc28f5c28f5c3p-1`, `-0x1.136694898f606p-3`, `-0x1.df1b8aa00192ap+0`, `-0x1.fd70a3d70a3d7p-1`, `-0x1.fd70a3d70a3d7p-1`, `-0x1.fd70a3d70a3d7p-1`, `-0x1.8f61672324c83p-1`, `-0x1.8f61672324c83p-1`, `-0x1.fd70a3d70a3d7p-1`, `-0x1.fd70a3d70a3d7p-1`, `-0x1.fd70a3d70a3d7p-1`, `-0x1.752f7f498c3b1p-1`, `-0x1.3e9a6f826edabp-2`, `-0x1.84906034f3fd9p-1`, `-0x1.e693c03bc4d23p-2`, `-0x1.fd70a3d70a3d7p-1`, `-0x1.fd70a3d70a3d7p-1`, `-0x1.fd70a3d70a3d7p-1`, `-0x1.9dcad14a0a0f5p-1`, `-0x1.9dcad14a0a0f5p-1`, `-0x1.9dcad14a0a0f5p-1`, `-0x1.fd70a3d70a3d7p-1`, `-0x1.fd70a3d70a3d7p-1`, `-0x1.fd70a3d70a3d7p-1`, `-0x1.5cd8dbcec84f9p-2`, `-0x1.20780fdc1615fp-1`, `-0x1.3bc7f77af6406p-3`, `-0x1.a0807357e670ep-4`, `0x1.cf7867f0aa21ap+0`, `-0x1.202c6ef3d3a1dp-1`, `0x1.e116ebd4cfd09p-2`, `0x1.724894c447c31p-2`, `-0x1.fd70a3d70a3d7p-1`, `-0x1.fd70a3d70a3d7p-1`, `-0x1.fd70a3d70a3d7p-1`, `0x1.0a6fc58ab92c0p-4`, `0x1.ed7add15f02c5p-2`, `-0x1.a8130164840e1p-1`, `-0x1.69014b599aa61p-6`, `-0x1.fd70a3d70a3d7p-1`, `0x1.49fe86833c600p+0`, `0x1.df69878316a05p-5`, `0x1.141d3aa369fcfp+0`, `0x1.345a1cac08312p-2`.

Certificate of record at $M = 7{,}680$:

U	$-205622732434150043/200000000000000000$
L_{free}	$-50720023282629301/50000000000000000$
L_{ext}	$-1011927969250662981/1000000000000000000$
margin	$2742639303632839/205622732434150043$

P7, case $(1/30, 4/15)$**: 49 constants.**

class	value	class	value	class	value
$(0,3,2)$	-0.399106	$(0,3,3)$	-0.930000	$(1,2,2)$	-1.606746
$(1,2,3)$	-1.849730	$(2,1,3)$	-0.968353	$(2,2,1)$	-0.995000
$(2,2,2)$	-0.995000	$(2,2,3)$	-0.995000	$(2,3,1)$	-0.677161
$(2,3,2)$	-0.677161	$(3,1,1)$	-0.995000	$(3,1,2)$	-0.995000
$(3,1,3)$	-0.995000	$(3,2,1)$	-0.699634	$(3,2,2)$	-0.579943
$(3,2,3)$	-0.699634	$(4,1,1)$	-0.995000	$(4,1,2)$	-0.995000
$(4,1,3)$	-0.995000	$(4,2,1)$	-0.812334	$(4,2,2)$	-0.812334
$(4,2,3)$	-0.812334	$(4,3,1)$	0.842426	$(5,0,0)$	-0.995000
$(5,1,1)$	-0.995000	$(5,1,2)$	-0.995000	$(5,1,3)$	-0.920991
$(5,2,1)$	-0.995000	$(5,2,2)$	-0.298363	$(5,2,3)$	-0.995000
$(6,1,1)$	-0.428380	$(6,1,2)$	-0.583524	$(6,1,3)$	-0.028459
$(6,2,1)$	-0.995000	$(6,2,2)$	1.573511	$(6,2,3)$	-0.849436
$(7,0,0)$	-0.995000	$(7,1,1)$	-0.905487	$(7,1,2)$	-0.995000
$(7,1,3)$	1.539724	$(7,2,1)$	-0.971748	$(8,0,0)$	-0.262985
$(8,1,1)$	-0.995000	$(8,1,2)$	-0.995000	$(8,1,3)$	1.357521
$(9,0,0)$	-0.988371	$(10,0,0)$	-0.917609	$(11,0,0)$	0.865323
$(12,0,0)$	0.318982				

Exact values, in the order of the table: `-0x1.98af3e468cac5p-2`, `-0x1.dc28f5c28f5c3p-1`, `-0x1.9b53b4b2fa93bp+0`, `-0x1.d987e7c06e19cp+0`, `-0x1.efcbf6e3f78bcp-1`, `-0x1.fd70a3d70a3d7p-1`, `-0x1.fd70a3d70a3d7p-1`, `-0x1.fd70a3d70a3d7p-1`, `-0x1.5ab4d8ba40d91p-1`, `-0x1.5ab4d8ba40d91p-1`, `-0x1.fd70a3d70a3d7p-1`, `-0x1.fd70a3d70a3d7p-1`, `-0x1.fd70a3d70a3d7p-1`, `-0x1.66366d7a56de3p-1`, `-0x1.28ee49f51697fp-1`, `-0x1.66366d7a56de3p-1`, `-0x1.fd70a3d70a3d7p-1`, `-0x1.fd70a3d70a3d7p-1`, `-0x1.fd70a3d70a3d7p-1`, `-0x1.9fea3df6db941p-1`, `-0x1.9fea3df6db941p-1`, `-0x1.9fea3df6db941p-1`, `0x1.af5275ee99a63p-1`, `-0x1.fd70a3d70a3d7p-1`, `-0x1.fd70a3d70a3d7p-1`, `-0x1.fd70a3d70a3d7p-1`, `-0x1.d78c21e1d2179p-1`, `-0x1.fd70a3d70a3d7p-1`, `-0x1.318611fd5885dp-2`, `-0x1.fd70a3d70a3d7p-1`, `-0x1.b6a93f290abb4p-2`, `-0x1.2ac3a860dcb9bp-1`, `-0x1.d245b291b823dp-6`, `-0x1.fd70a3d70a3d7p-1`, `0x1.92d19dece5711p+0`, `-0x1.b2e9468017119p-1`, `-0x1.fd70a3d70a3d7p-1`, `-0x1.cf9bfdf7e8039p-1`, `-0x1.fd70a3d70a3d7p-1`, `0x1.8a2b5a20ddc62p+0`, `-0x1.f188f42fe8251p-1`, `-0x1.0d4bf0995aaf8p-2`, `-0x1.fd70a3d70a3d7p-1`,

`-0x1.fd70a3d70a3d7p-1`, `0x1.5b867f0aa21a7p+0`, `-0x1.fa0bc382a12f9p-1`, `-0x1.d5d0d8cb07d0bp-1`, `0x1.bb0b9dc2f405fp-1`, `0x1.46a337a80cf9ep-2`.

Certificate of record at $M = 23{,}040$:

U	$-204198366746515443/200000000000000000$
L_{free}	$-253694362299242313/250000000000000000$
L_{ext}	$-101553198415196703/100000000000000000$
margin	$363989972040679/68066122248838481$

Appendix B. Kernels and rows

Unless a second argument is written, the functions below are functions of one variable on $[0,1]$, a total size; ρ and M_j take two (a subset sum ς and a total τ), and XS_k and $\Lambda_k^{(j)}$ carry a second argument as shown. Put $s(u) = 1/u$ on $[\eta, \alpha']$ and $r(u) = 1/u$ on $(\alpha', 1/2)$, both zero elsewhere, and

$$S_j = s^{*j}/j!, \qquad R_k = r^{*k}/k!, \qquad S_0 = R_0 = \delta_0.$$

$R_{k,b}$ is the part of R_k whose largest component lies in bin b, and $D_{j,k,b} = S_j * R_{k,b}$. "tr" marks multiplication by $\mathbf{1}_{[1/2,1]}$.

Fragmentation kernels. A large component $a \ge 1/2$ of an extended vector is fragmented as in Lemma 2.1. Summing Lemma 2.1 over the fragmentations of a into j small and k rough fragments, $j+k \ge 2$, with B running over all sub-collections of the fragments (b_s small and b_r rough members), gives

$$\Lambda_{j,k}(a) = \sum_{b_s=0}^{j} \sum_{b_r=0}^{k} (-1)^{j+k+1+b_s+b_r} \big((S_{b_s} * R_{b_r}) \mathbf{1}_{[1/2,1]} * S_{j-b_s} * R_{k-b_r} \big)(a), \qquad a \ge 1/2;$$

the boundary term of Lemma 2.1 at $a = 1$ integrates to zero. For a fragment type $y = (j_y, k_y, b_y)$ of the large component, Λ_y is the part of Λ_{j_y,k_y} whose largest rough fragment lies in bin b_y ($\Lambda_y = \Lambda_{j_y,0}$ when $k_y = 0$), and $c_y := c[j_y, k_y, b_y]$. These are the kernels of the constant part and of $f(1)$ below. The terms with $b_r = 0$, sub-collections made only of small fragments, are nonzero as soon as the small fragments can total at least $1/2$, and for $k = 0$ they are the whole kernel.

The forced part meets only fragmentations whose small fragments total at most $\alpha' < 1/2$. There no sub-collection without a rough fragment reaches $1/2$, and the rough fragments can be resolved by number: for $k \ge 1$ rough fragments and a threshold $\vartheta > 0$, put

$$XS_k(t; \vartheta) = (-1)^k \sum_{b_r=1}^{k} (-1)^{b_r} \big(R_{b_r} \mathbf{1}_{[\vartheta,\infty)} * R_{k-b_r} \big)(t).$$

With no small fragment, $\Lambda_k^0(a) = -XS_k(a; 1/2) = \Lambda_{0,k}(a)$ for $a \ge 1/2$; $\Lambda_1^0 \equiv 0$ there, since a fragmentation has at least two pieces. With j small fragments of total $\tau \le \alpha'$,

$$\Lambda_k^{(j)}(a; \tau) = (-1)^{j-1} \int M_j(\varsigma, \tau)\, XS_k(a - \tau; 1/2 - \varsigma)\, d\varsigma, \qquad M_j = \rho^{*j}/j!,$$

where $\rho(\varsigma, \tau) = s(\tau)[\delta(\varsigma) - \delta(\varsigma - \tau)]$ is the signed measure of one small fragment on (subset sum ς, total τ), and the convolution power is taken in both variables; equivalently $M_j(\varsigma, \tau) = \sum_{b=0}^{j} (-1)^b S_b(\varsigma) S_{j-b}(\tau - \varsigma)$. The kernels with bins are resolved by the bin of the largest rough fragment. When every rough part exceeds $1/4$, $\Lambda_2^0 = -R_2$ and $\Lambda_3^0 = 2R_3$ on $[1/2, 1]$.

The rows. Consider a modulus with j_D small prime factors of total σ ($j_D = 0$, $\sigma = 0$ for $d = 1$), and put $t = 1 - \sigma$. Every vector in the row carries the sign $(-1)^{\dim}$ of its whole dimension.

- The constant part is

$$\sum_{(j',k',b')} (-1)^{j_D+j'+k'} c[j_D+j',k',b']\, D_{j',k',b'}(t) + \sum_{w} (-1)^{1+j_D+j_w+k_w} \sum_{y} c[j_D+j_w+j_y,\ k_w+k_y,\ \max(b_w,b_y)]\, (D_w * \Lambda_y \,\mathrm{tr})(t),$$

  where $w=(j_w,k_w,b_w)$ runs over the free remainders, the empty one included ($D_\emptyset=\delta_0$) except in T_0 (below), and y over the fragment types of the large component.
- The forced part, for $h_J(\tau)$ with $J=j_D+d_j$ and $u=\tau-\sigma$, has free completions and extended completions. The free completion with $d_j=0$ and k rough parts gives the diagonal term $(-1)^{j_D+k}w_kR_k(t)$; free completions with $d_j\ge 1$ give $(-1)^{J+k}w_kS_{d_j}(u)R_k(1-\tau)$. Extended completions, with $j_w+j_y=d_j$, $k_w+k_y=k$, $k_y\ge 1$ and sign $(-1)^{1+j_D+j_w+k_w}$, give four kinds of terms: $(R_{k_w}*\Lambda^0_{k_y}\mathrm{tr})(t)$, $S_{j_w}(u)(R_{k_w}*\Lambda^0_{k_y}\mathrm{tr})(1-\tau)$, $(R_{k_w}*\Lambda^{(j_y)}_{k_y}(\cdot;u)\mathrm{tr})(t)$, and $\int S_{j_w}(v-\sigma)\,(R_{k_w}*\Lambda^{(j_y)}_{k_y}(\cdot;\tau-v)\mathrm{tr})(1-v)\,dv$, each with weight w_k.
- The diagonal is the sum of the terms with $d_j=0$: $B_{j_D}(\sigma)=\sum_k w_k\big[(-1)^{j_D+k}R_k(t)+\sum_{k_w+k_y=k,\ k_y\ge1}(-1)^{1+j_D+k_w}(R_{k_w}*\Lambda^0_{k_y}\mathrm{tr})(t)\big]$. Lemma 3.1 evaluates it for the weighted profile.
- T_0 is the row with $j_D=0$. Every forced term in it has $d_j=J\ge1$; on extended vectors these terms come with every nonempty leftover vector, including those with no rough part. The prime vector, whose single component 1 is the large component with empty remainder, contributes $-f(1)$; this term is written separately and counted once, so in T_0 the sum over w runs over nonempty remainders only.
- $f(1)=\sum_y c_y\Lambda_y(1)+\sum_J\int h_J(\tau)\,\Lambda^{(J)}_k(1;\tau)\,d\tau$, summed over the profile with weights w_k.

The untwisted rows, the same sums without the signs, vanish identically for any free data, since every row has $\sigma\le\alpha'<1/2$, where (6.2) applies (§2); the certificates check this, to the extent stated in §5, as a test of the machinery.

Appendix C. Proofs of the enclosure lemmas

We use the notation of §5.4: M is the resolution, $x_n=n/M$, $K_n=[x_n,x_{n+1}]$ and K_n° is its interior. A cellwise affine function F has end values $a_n=F(x_n+)$ and $b_n=F(x_{n+1}-)$ on K_n and vanishes outside $[0,1)$. Its distributional derivative is

$$F'=\sum_n M\,\Delta a_n\,\mathbf{1}_{K_n}+\sum_n J_n^F\,\delta_{x_n},\qquad \Delta a_n=b_n-a_n,\quad J_n^F=a_n-b_{n-1},$$

with $b_{-1}=0$, and $J_M^F=-b_{M-1}$ the jump at 1. We use the interpolation remainder: if φ is continuous on K_k and twice differentiable on K_k°, and ℓ is its chord on K_k, then for every $x\in K_k$ there is $\xi\in K_k^\circ$ with

$$\varphi(x)-\ell(x)=-\tfrac12(x-x_k)(x_{k+1}-x)\,\varphi''(\xi),\qquad 0\le(x-x_k)(x_{k+1}-x)\le\frac{1}{4M^2}. \tag{7}$$

It follows from Rolle's theorem applied twice to $t\mapsto\varphi(t)-\ell(t)-c\,(t-x_k)(t-x_{k+1})$, with c chosen so that this function vanishes at x.

Proof of Lemma 5.1. On K_i, $1/u$ is decreasing, which gives the first pair of bounds, and convex, so it lies below its chord ℓ_i. Its second derivative is $2/u^3\le 2M^3/i^3$ on K_i, so by (7) $1/u\ge\ell_i(u)-\frac{1}{8M^2}\cdot\frac{2M^3}{i^3}=\ell_i(u)-\frac{M}{4i^3}$. □

Proof of Lemma 5.2. F and G are bounded with bounded support, so $s = F * G$ is continuous. (a) Only pairs of cells with $n + m = k - 1$ contribute to $s(x_k)$. Put $y = x_n + \theta/M$ with $\theta \in [0, 1]$; then $x_k - y$ lies in K_m at relative position $1 - \theta$, and

$$\int_{K_n} F(y)\, G(x_k - y)\, dy = \frac{1}{M} \int_0^1 \big(a_n + \Delta a_n \theta\big)\big(d_m - \Delta c_m \theta\big)\, d\theta = \frac{1}{6M}\big(a_n c_m + b_n d_m + 2a_n d_m + 2b_n c_m\big).$$

(b) We have $s'' = F' * G'$. The term $M^2 \Delta a_n \Delta c_m\, \mathbf{1}_{K_n} * \mathbf{1}_{K_m}$ is a hat function: it is affine on K_{n+m}, rising from 0 to $M \Delta a_n \Delta c_m$, and affine on K_{n+m+1}, falling back to 0. The term $M J_n^F \Delta c_m\, \delta_{x_n} * \mathbf{1}_{K_m} = M J_n^F \Delta c_m\, \mathbf{1}_{K_{n+m}}$ is constant on K_{n+m}, and likewise $M \Delta a_n J_m^G\, \mathbf{1}_{K_{n+m}}$. The terms $J_n^F J_m^G\, \delta_{x_{n+m}}$ are carried by grid points. So s'' is affine on K_k°; at x_k+ the hats with $n + m = k - 1$ contribute their peaks and those with $n + m = k$ contribute 0, and at $x_{k+1}-$ the reverse. This gives MX_k^- and MX_k^+. Jumps at 1 act only on cells beyond 1. (c) s is continuous on K_k and, by (b), twice differentiable on K_k° with s'' affine there, so the extreme values of s'' on K_k° are its one-sided end values. By (7), $s - s_k \geq -\frac{1}{8M^2} \max(0, MX_k^-, MX_k^+)$ and $s - s_k \leq \frac{1}{8M^2} \max(0, -MX_k^-, -MX_k^+)$. For the correlation $r(x) = \int F(u)\, G(u - x)\, du$ the same computation, with $u = x_n + \theta/M$ and $u - x_k$ at relative position θ in K_{n-k}, gives the stated knot values, and two integrations by parts give $r'' = -\int F'(u)\, G'(u - x)\, du$; for $0 \leq x < 1$ the jump of F at 1 meets G at $1 - x \in (0, 1]$, so it is included. Finally, if $0 \leq F^- \leq f \leq F^+$ and $0 \leq G^- \leq g \leq G^+$ on open cells, then $F^- * G^- \leq f * g \leq F^+ * G^+$ everywhere, since the integrands are ordered almost everywhere; (c) applied to the two outer convolutions encloses $f * g$, and outward rounding of the knot values and of the slack keeps it enclosed. □

Proof of Lemma 5.3. Fix an open cell. If both end values of F^- are ≥ 0, then $F^- \geq 0$ on the cell, so $f^+ = f \geq F^-$; otherwise $f^+ \geq 0$ is the claimed lower bound. Next, $f^+ = \max(f, 0) \leq \max(F^+, 0)$, and $\max(F^+, 0)$ is convex on the cell, being the maximum of two affine functions, so it lies below its chord. The statements for $f^- = (-f)^+$ follow from these, applied to $-F^+ \leq -f \leq -F^-$. Since $f = f^+ - f^-$ with $f^\pm \geq 0$, a product or convolution of signed functions is a signed sum of four products or convolutions of nonnegative functions, each enclosed by Lemma 5.2 or 5.4; its lower bound uses the lower enclosures of the terms with sign + and the upper enclosures of the terms with sign −, and its upper bound the reverse. □

Proof of Lemma 5.4. Fix $\tau \in [0, 1]$. (a) The map $(p, q) \mapsto pq$ is affine in each variable, so on the box $[p^-(\tau), p^+(\tau)] \times [q^-(\tau), q^+(\tau)]$ its extreme values are taken at corners. Each corner function $r = p^\pm q^\pm$ is quadratic in τ with $r'' = 2\Delta p^\pm \Delta q^\pm$, so by (7) (in the variable τ, with cell length 1) it lies within $|\Delta p^\pm \Delta q^\pm|/4$ of its chord. Hence $pq \leq \max_r(\ell_r + \sigma_r)$, where ℓ_r is the chord of r and σ_r its slack. The right side is a maximum of affine functions of τ, hence convex, hence below its own chord, which is the chord through the end values $\max_r(r(0) + \sigma_r)$ and $\max_r(r(1) + \sigma_r)$. The lower bound is symmetric: a minimum of affine functions is concave, so it lies above its chord. (b) If $q^- > 0$ throughout the closed cell or $q^+ < 0$ throughout it, then on the box the map $(p, q) \mapsto -p/q$ is monotone in p and in q, so again its extreme values are at corners. For a corner $r = -p/q$ with p, q affine in τ, $r' = -(\Delta p\, q - p\, \Delta q)/q^2$ and $\Delta p\, q - p\, \Delta q = \Delta p\, q_0 - p_0\, \Delta q$ is constant, so $r'' = 2\Delta q\, (\Delta p\, q_0 - p_0\, \Delta q)/q^3$. As q is affine and, by the hypothesis, does not vanish on the closed cell, $|q| \geq \min(|q_0|, |q_1|)$ there, which gives the stated bound, and the argument of (a) applies with the slack $\max |r''|/8$. □

Proof of Lemma 5.5. (a) By induction downwards from $j = m^*$: the coupling terms $C_{j,i}[h_{j+i}]$ with $i \geq 1$ involve only levels already enclosed, and they are built from convolutions, correlations, cell products and block bounds of enclosed functions (Lemmas 5.1 to 5.4 and 5.6). So the numerator of $h_j = -(A_j \cdot c + \sum_{i \geq 1} C_{j,i}[h_{j+i}])/B_j$ is enclosed on every cell, and Lemma 5.4(b) encloses the quotient, because the enclosure of B_j excludes zero on the closed cell. (b) Write $\langle u, v \rangle = \int uv$ over the forcing range and let $C_{j,i}^*$ be the transpose of $C_{j,i}$, so that $\langle \lambda, C_{j,i}[g] \rangle = \langle C_{j,i}^*[\lambda], g \rangle$ (Fubini; all

functions are bounded with bounded support). The recursion says $B_m\lambda_m + \sum_{j<m} C^*_{j,m-j}[\lambda_j] = \kappa_m$ for every m, and the forcing says $B_j h_j + \sum_{i\geq 1} C_{j,i}[h_{j+i}] = -A_j \cdot c$ for every j. Hence

$$\begin{aligned}\sum_m \langle \kappa_m, h_m\rangle &= \sum_m \langle B_m\lambda_m, h_m\rangle + \sum_{j<m} \langle C^*_{j,m-j}[\lambda_j], h_m\rangle = \sum_j \Big\langle \lambda_j,\ B_j h_j + \sum_{i\geq 1} C_{j,i}[h_{j+i}]\Big\rangle \\ &= -\sum_j \langle \lambda_j, A_j \cdot c\rangle.\end{aligned}$$

The recursion runs upwards from $m = 1$, and each step is enclosed as in (a), with the enclosure of B_m excluding zero on every closed cell; so the exact λ_m are enclosed, and then so is the right side. The left side is also enclosed directly from the enclosures of h_m. Both enclosures contain the exact value, so their intersection does. □

Proof of Lemma 5.6. (a) For $\tau_0 \leq \tau \leq \tau_1$, $\mathbf{1}_{[1/2-\tau_0,1/2)} \leq \mathbf{1}_{[1/2-\tau,1/2)} \leq \mathbf{1}_{[1/2-\tau_1,1/2)}$, and integrating against a nonnegative factor preserves the order. (b) The difference $X_\tau(x) - X_{\tau_0}(x)$ is the integral of $r(1/2-t)\, g(x-1/2+t)$ over $t \in [\tau_0, \tau]$, and $r \geq 0$; replacing g by its minimum and its maximum over the arguments that occur gives the bounds. When the correction is integrated against a signed function H, the bounds are applied separately to the integrals against H^+ and H^- (Lemma 5.3), since g depends on the integration variables. (c) is immediate. Under the additional hypotheses stated in the lemma, bounded integrated density proves the rate in (a), and the Lipschitz range bound proves it in (c). In (b), bounded r gives $\rho(\tau) = O(\tau_1 - \tau_0)$ and the Lipschitz bound gives variation $O(\tau_1 - \tau_0 + 1/M)$ on the window. Multiplying proves the stated additional block loss. These rate arguments are not used on windows crossing a jump. □

Proof of Lemma 5.7. (a) On the cell, $a\,\Phi(a) \geq a\,L(a) =: q(a)$, as $a > 0$, and q is quadratic with $q'' = 2M(L_1 - L_0)$. If $L_1 \leq L_0$, q is concave and its minimum is at an end of the cell. Otherwise (7) gives $q \geq \ell - \frac{1}{8M^2} \cdot 2M(L_1 - L_0)$, and the chord ℓ of q is at least the smaller of its end values $x_n L_0$ and $x_{n+1} L_1$. (b) For fixed a, minimising over t and over the interval for f_g uses its four corners. Replacing f_c by L_c gives the stated minimum of affine functions, hence a concave function above its end chord. (c) Pass to the limit from within a cell. The assertion about the assigned point value then follows from the stated convention; it does not follow without that convention. The remaining collapsed strata are checked directly from the class constants and the support of the forced terms. □

Proof of Lemma 5.8. (a) Round to nearest returns a binary64 number y with $|y - z| \leq \frac{1}{2}$ulp for the exact result z, where the unit in the last place is that of the binade in which the rounding takes place, unless the operation overflows. So z lies between the neighbours of y below and above. For (b), expand each computed term and each addition along its path to the final sum as multiplication by factors $1 + \delta$, with $|\delta| \leq u$. A path has at most $k - 1 + r$ factors, and the standard product bound [10, Lemma 3.1] gives $|\prod(1+\delta) - 1| \leq \gamma_{k-1+r}$ when $(k-1+r)u < 1$. Taking absolute values and summing proves the claim. For additions of binary64 inputs an exact subnormal result is representable, so such an addition is exact; the relative-error hypothesis holds for products and quotients with normal results. Source-specific removal of true-zero terms in the raw nonnegative accumulators is described after the lemma. Individually widened operations instead use part (a), including for subnormal results. (c) Write $x = k\,2^e$ with integers $|k| < 2^{53}$ and $e \geq -1074$. Then $x\,2^S = k\,2^{e+S}$ with $e + S \geq -1074$ and $|x\,2^S| < 2^{52}$, so it is a binary64 number and the multiplication is exact. The integers $\lfloor x2^S \rfloor$ and $\lceil x2^S \rceil$ have absolute value at most 2^{52}, so they are exactly representable, and they bound $x2^S$ below and above. Convolutions of such integer images are exact integer polynomial products. □

References

[1] E. Bombieri, *On the large sieve*, Mathematika **12** (1965), 201–225.

[2] E. Bombieri, *The asymptotic sieve*, Rend. Accad. Naz. XL (5) **1/2** (1975/76), 243–269 (1977).
[3] K. Ford, *On Bombieri's asymptotic sieve*, Trans. Amer. Math. Soc. **357** (2005), 1663–1674.
[4] K. Ford, *Prime producing sieves* (joint work with J. Maynard), in: Analytic Number Theory, Mathematisches Forschungsinstitut Oberwolfach, Report No. 51/2025, 2727–2729, doi:10.4171/OWR/2025/51.
[5] K. Ford and J. Maynard, *On the theory of prime-producing sieves*, arXiv:2407.14368v1 (2024).
[6] J. Friedlander and H. Iwaniec, *Asymptotic sieve for primes*, Ann. of Math. (2) **148** (1998), no. 3, 1041–1065.
[7] J. Friedlander and H. Iwaniec, *Opera de Cribro*, Amer. Math. Soc. Colloquium Publications **57**, Amer. Math. Soc., Providence, RI, 2010.
[8] J. Friedlander and H. Iwaniec, *Coordinate distribution of Gaussian primes*, J. Eur. Math. Soc. **24** (2022), 737–772.
[9] J. Guo, *The Möbius function at $N - p$*, preprint, version 2, 23 October 2026, 65 pp. Software Heritage archive, revision `swh:1:rev:41956ce2ad80d301a0fa8929b518603be4881196`, file `mobius-goldbach-shift/paper.pdf`, content `swh:1:cnt:481af3a5d4f923850ed07848b931c23106512d8b`; `https://archive.softwareheritage.org/swh:1:cnt:481af3a5d4f923850ed07848b931c23106512d8b/`.
[10] N. J. Higham, *Accuracy and Stability of Numerical Algorithms*, 2nd ed., SIAM, Philadelphia, 2002.
[11] J.-J. Huang and H. Li, *On the connection between the Goldbach conjecture and the Elliott–Halberstam conjecture*, in: M. B. Nathanson (ed.), Combinatorial and Additive Number Theory IV, Springer Proc. Math. Stat. **347**, Springer, Cham, 2021, 323–346; arXiv:2005.03811v2.
[12] H. Iwaniec, *Prime numbers and L-functions*, in: Proceedings of the International Congress of Mathematicians (Madrid, 2006), Vol. I, European Mathematical Society, Zürich, 2007, 279–306.
[13] H. Iwaniec and E. Kowalski, *Analytic Number Theory*, Amer. Math. Soc. Colloquium Publications **53**, Amer. Math. Soc., Providence, RI, 2004.
[14] G. Lü, Z. Wang and J. Wu, *On the Hardy–Littlewood–Chowla conjecture*, arXiv:2610.20242v1 (2026).
[15] R. Moya, *A moving-cut correction to the Huang–Li conditional Goldbach argument: consequences for diagonal Möbius-twisted Elliott–Halberstam hypotheses*, preprint, Zenodo, 2026, doi:10.5281/zenodo.21939906.
[16] M. R. Murty and A. Vatwani, *Twin primes and the parity problem*, J. Number Theory **180** (2017), 643–659.
[17] A. Selberg, *The general sieve-method and its place in prime number theory*, in: Proceedings of the International Congress of Mathematicians (Cambridge, Mass., 1950), Vol. 1, Amer. Math. Soc., Providence, R.I., 1952, 286–292.
[18] A. Vatwani, *Higher rank sieves and applications*, PhD thesis, Queen's University, 2016.
[19] A. I. Vinogradov, *The density hypothesis for Dirichlet L-series*, Izv. Akad. Nauk SSSR Ser. Mat. **29** (1965), 903–934 (Russian).