IMPROVED BOUNDS FOR THE PRIMITIVE-SET SATURATION GAME (ERDŐS PROBLEM 872)

OM BUDDHDEV

Abstract. Erdős posed the divisibility antichain saturation game on $\{2, 3, \ldots, n\}$ in 1992 [Erd92], catalogued as Problem #872 in Bloom's list [Blo26]. Two players alternately choose integers, maintaining the property that the chosen set is primitive (an antichain under divisibility); Prolonger moves first to maximize the final count, Shortener moves to minimize it. The central question is whether the game value $L(n)$ is linear in n or sublinear. We prove

$$\left(\frac{1}{8} - o(1)\right) \frac{n \log \log n}{\log n} \ \le\ L(n) \ \le\ \left(\frac{\mathcal{W}_4}{2} + o(1)\right) n \ <\ 0.19\, n,$$

where $\mathcal{W}_4$ is an explicit fourth-order Bonferroni constant arising from the lower logarithmic profile of Shortener's legal-prime prefix. The lower bound follows from a two-phase Prolonger strategy and a bipartite maximum-degree capture lemma; the upper bound uses a monotone envelope for the legal-prime lower profile, a queued prime-rounding bridge, and a fourth-order Bonferroni inequality on the odd-part compression of the terminal antichain. We also prove a polynomial lower bound for the weighted shield functional, an exact $5n/24 + O(1)$ first-hit cover theorem for the lower half, and intermediate asymptotic upper-bound constants 13/36 and 5/16. Conditional on a restricted finite safe-edge hypothesis—whose general form we refute in the paper—we also obtain a stronger lower bound of order $n(\log \log n)^2/ \log n$. The paper also records three proof-class obstructions for related proof methods. Shield reduction, the exact $5n/24$ cover identities, the 13/36 upper-bound core, and the $\mathcal{W}_4/2 < 0.19$ endgame reduction have zero-sorry Lean 4 artifacts; the formalization status of remaining results is documented in Appendix C.

1. Introduction

Let $n \ge 2$. The divisibility antichain saturation game is played on

$$\{2, 3, \ldots, n\}.$$

Two players alternately choose previously unchosen integers, and the chosen set must remain an antichain under divisibility. Thus no chosen integer may divide another chosen integer. The game stops when the current antichain is maximal. Prolonger moves first and tries to maximize the final number of chosen integers; Shortener tries to minimize it. We denote the value of the game by $L(n)$.

This game is one of the finite saturation games arising from Erdős's questions on primitive sets [Erd92]; saturation games of this kind are part of a broader extremal-game tradition [FS92, BHW16]. A primitive set is precisely an antichain in the divisibility poset, and primitive sets have a long parallel history in multiplicative number theory [CE90, Lic23, MP10, Kuc24]. We write

$$L_n := \{m \in \mathbb{Z} :\ 2 \le m \le \lfloor n/2 \rfloor\}, \qquad U_n := \{m \in \mathbb{Z} :\ n/2 < m \le n\}.$$

The upper half U_n is the canonical maximum antichain, so the trivial scale is linear. The central question is whether optimal play forces a positive fraction of this linear scale, or whether Shortener can force $o(n)$ moves.

Date: April 21, 2026.

2020 *Mathematics Subject Classification.* Primary 91A46; Secondary 05D05, 11N25.

Key words and phrases. Saturation game, primitive set, divisibility poset, antichain, sieve methods, formal verification, Erdős problem 872.

1.1. **Prior work.** Erdős posed the problem in 1992 [Erd92], and it is catalogued as problem #872 in Bloom's list [Blo26]. No non-trivial upper bound appeared in the public record until early 2026. The game is a number-theoretic analogue of the triangle-free saturation game of Hajnal: for that game, Füredi and Seress [FS92] proved a lower bound of order $n \log n$, while Biró, Horn, and Wildstrom [BHW16] proved an upper bound below $0.215n^2$.

The ambient combinatorial setting interacts with the literature on primitive sets and divisibility antichains. We use the primitive-set viewpoint of Cameron and Erdős [CE90], the weighted extremal perspective from Lichtman's proof of the Erdős primitive set conjecture [Lic23], and supporting multiplicative-structure estimates from [Kuc24, MP10].

The public upper-bound record for the present game developed rapidly in early 2026. To the author's knowledge, the first linear upper bound,

$$L(n) \leq \left(\tfrac{23}{48} + o(1)\right) n,$$

was given by Price and GPT-5.2 Pro, who also introduced the terminology *Prolonger* and *Shortener* [Pri26]. Subsequent refinements recorded in the public forum thread by Adenwalla, StijnC, natso26, Xiao_Hu, Desmond Weisenberg, and others improved this to

$$L(n) \leq \left(\tfrac{419}{1008} + o(1)\right) n \leq 0.416n$$

[Erd26]. The results in this paper substantially strengthen that record and establish the first unconditional lower bound of order strictly larger than the trivial $n/\log n$ baseline (see Section 4).

1.2. **Main results.** Our main result is a new linear upper bound with an explicit constant below 0.19. Let

$$\rho(u) := \frac{1}{(\lfloor 1/u \rfloor + 1)u} \qquad (0 < u \leq 1),$$

and define

$$J_r := \frac{1}{r!} \int_{\substack{u_1 + \cdots + u_r \leq 1 \\ 0 < u_i \leq 1}} \prod_{i=1}^{r} \rho(u_i)\, du_1 \cdots du_r \qquad (r \geq 1).$$

Set

$$\mathcal{W}_4 := 1 - J_1 + J_2 - J_3 + J_4.$$

The certified numerical bound in Proposition B.1 is

$$\frac{\mathcal{W}_4}{2} \leq 0.1897123371 < 0.19.$$

Theorem 1.1 (Main upper bound)**.** *Shortener has a strategy forcing*

$$L(n) \leq \left(\frac{\mathcal{W}_4}{2} + o(1)\right) n < 0.19n.$$

The proof of Theorem 1.1 is given in Section 7. Shortener plays a long prefix of the smallest legal odd primes. A one-sided lower-profile estimate controls which primes can be kept legal on the logarithmic scale $p = n^u$. We build a monotone envelope for this profile, invert it to obtain a comparison sequence, round the comparison sequence to genuine primes without losing the first four factorial moments, and then apply a fourth-order Bonferroni inequality to the odd-part compression of the final antichain.

Theorem 1.2 (Unconditional fan lower bound)**.** *Prolonger has a strategy forcing*

$$L(n) \geq \left(\frac{1}{8} - o(1)\right) \frac{n \log \log n}{\log n}.$$

The proof of Theorem 1.2 is given in Theorem 4.4. Prolonger first works through the small odd primes $a \le n^\delta$, making them unavailable by legal upper-half multiples unless Shortener chooses the singleton first. A harmonic charging argument shows that the small primes made unavailable by proper multiples still have reciprocal mass at least half of the total small-prime mass. The remaining upper-half targets ab form a fan graph, and a maximum-degree right-capture lemma forces at least half of the surviving fan edges to become actual game moves.

The lower-bound side also contains a weighted obstruction theorem for the standard shield-reduction approach. For $x \in L_n$, let

$$M_n(x) := \{u \in U_n : \ x \mid u\}, \qquad w_n(x) := |M_n(x)| - 1.$$

For $P \subseteq U_n$, put

$$\mathcal{L}_n(P) := \{x \in L_n : \ x \nmid u \text{ for every } u \in P\},$$

and let

$$\beta_n(P) := \max\left\{\sum_{x\in B} w_n(x) : B \subseteq \mathcal{L}_n(P) \text{ is a divisibility antichain}\right\}.$$

Theorem 1.3 (Theorem A: Polynomial shield-weight lower bound)**.** *Fix $0 < \alpha < 1$. Uniformly for every $P \subseteq U_n$ with $|P| \le n^\alpha$,*

$$\beta_n(P) \ge \left(\frac{1}{2}\log\frac{1}{\alpha} - o(1)\right) n.$$

Equivalently, if

$$k_c^\star(n) := \min\{|P| : \ P \subseteq U_n, \ \beta_n(P) \le cn\},$$

then for every fixed $c > 0$,

$$k_c^\star(n) \ge n^{e^{-2c} - o(1)}.$$

The next theorem explains why this weighted quantity is the right one for shield-prefix arguments.

Theorem 1.4 (Shield reduction framework)**.** *Let A be any terminal antichain and let $P \subseteq A \cap U_n$. Then*

$$|A| \ge |U_n| - \beta_n(P).$$

Corollary 1.5 (Shield-prefix obstruction)**.** *Any proof which uses a prefix $P \subseteq U_n$ and the shield reduction framework to force $|A| \ge (1/2 - c)n + o(n)$ must have*

$$|P| \ge n^{e^{-2c} - o(1)}.$$

The proof of the reduction is elementary and is given in Section 3; Theorem 1.3 is proved in Section 4. These structural results show that a short upper-half shield prefix cannot certify a linear lower bound through the weighted shadow method.

We also determine exactly the size of the first-hit cover of the lower half by upper-half multiples. Define

$$\tau(n) := \min\{|C| : \ C \subseteq U_n, \ \forall x \in L_n \ \exists u \in C \text{ with } x \mid u\}.$$

Theorem 1.6 (Exact first-hit cover)**.**

$$\tau(n) = \frac{5}{24} n + O(1).$$

More precisely, the cover

$$H_n = \{u \in U_n : \ u \equiv 2 \pmod 4\} \cup \{u \in U_n : \ u > 2n/3, \ u \equiv 0 \pmod 4\}$$

has size $5n/24 + O(1)$, and there is a matching packing of lower-half witnesses of the same size up to $O(1)$.

Before the fourth-order argument, the same prime-prefix method gives two simpler upper bounds.

Theorem 1.7 (Intermediate upper bounds)**.** *Shortener has explicit strategies forcing*

$$L(n) \le \left(\frac{13}{36} + o(1)\right) n$$

and

$$L(n) \le \left(\frac{5}{16} + o(1)\right) n.$$

Finally, we include a conditional lower-bound theorem, not as a headline result but as a useful finite-combinatorial target. The natural general safe-edge hypothesis is false even inside the arithmetic residual hypergraphs, so the theorem is conditional on a narrower hypothesis for states generated by the activation and residual strategies constructed in Appendix A after the roadmap in Section 4. Choosing $\delta = 1/8$ in Theorem 4.7 gives, conditional on the restricted safe-edge hypothesis (Definition 4.5), an absolute constant $c > 0$ such that

$$L(n) \ge c\,\frac{n(\log\log n)^2}{\log n}$$

for all sufficiently large n. The finite graph, hypergraph, and arithmetic embedding cores of this theorem have Lean artifacts listed in Appendix C; the restricted safe-edge hypothesis and the asymptotic activation wrapper remain in prose and are stated explicitly rather than suppressed.

Several parts of the paper are accompanied by Lean formalizations; the exact artifact paths, theorem declaration names, repository snapshot commit, and Lean toolchain version are recorded in Appendix C. In summary, the Shield Reduction theorem, the structural identities underlying the exact 5/24 cover, and the 13/36 upper-bound core have zero-sorry Lean artifacts. Theorem A and the 5/16 upper bound have Lean artifacts whose remaining sorries are standard analytic-number-theory inputs or their assembly; the 5/16 artifact's remaining sorry also covers the game-tree wrapper realizing the prescribed prefix. The finite graph/hypergraph/embedding cores of the conditional T2 theorem have Lean artifacts, while the restricted safe-edge hypothesis and asymptotic wrapper remain prose. For the $0.19n$ theorem, the endgame reduction from moment convergence to the strict inequality is formalized; the moment-convergence hypotheses are the prose content of Section 7.

The paper is organized as follows. Section 2 fixes notation. Section 3 proves the shield reduction. Section 4 proves Theorem A, the unconditional fan lower bound, and states the conditional T2 theorem. Section 5 proves the 5/24 cover theorem. Section 6 proves the 13/36 and 5/16 upper bounds. Section 7 proves the main $0.19n$ bound. Section 8 records three obstructions to extending the present methods. Deferred proofs, numerical details, and formalization metadata appear in the appendices.

2. Notation and Setup

All divisibility statements in this paper are inside the finite board

$$V_n := \{2, 3, \dots, n\}.$$

A set $A \subseteq V_n$ is *primitive* if no two distinct elements of A are comparable under divisibility. Equivalently, A is an antichain in the divisibility poset. A primitive set is *maximal* if no element of $V_n \setminus A$ can be added while preserving primitiveness.

The game value $L(n)$ is the final size of the primitive set under optimal play when Prolonger moves first and maximizes the final size, while Shortener minimizes it.

We use the lower–upper split

$$L_n := \{m \in \mathbb{Z} :\ 2 \le m \le \lfloor n/2 \rfloor\}, \qquad U_n := \{m \in \mathbb{Z} :\ n/2 < m \le n\}.$$

The upper half U_n is an antichain, since two distinct elements of U_n cannot divide each other.

For $x \in L_n$, the upper shadow and shadow weight are

$$M_n(x) := \{u \in U_n : \ x \mid u\}, \qquad w_n(x) := |M_n(x)| - 1.$$

The subtraction by 1 is convenient because a lower-half move x adds one move but may remove $|M_n(x)|$ possible upper-half moves from a terminal position.

For $P \subseteq U_n$, define

$$\mathcal{L}_n(P) := \{x \in L_n : \ x \nmid u \text{ for every } u \in P\},$$

and

$$\beta_n(P) := \max\left\{\sum_{x\in B} w_n(x) : B \subseteq \mathcal{L}_n(P) \text{ is a divisibility antichain}\right\}.$$

The maximum is over a finite set, and is understood to be 0 when $\mathcal{L}_n(P)$ is empty.

We use standard arithmetic notation. The functions $\pi(x)$ and $\vartheta(x)$ denote the prime-counting function and Chebyshev's function, respectively. For an integer m, $P^-(m)$ and $P^+(m)$ denote the least and greatest prime factors when $m > 1$, $\Omega(m)$ denotes the number of prime factors counted with multiplicity, and $\omega(m)$ denotes the number of distinct prime factors. The symbol $\mathbb{P}$ denotes the set of primes. As usual, $O(\cdot)$, $o(\cdot)$, $\ll$, and $\gg$ are interpreted as $n \to \infty$ unless a different parameter is explicitly specified; subscripts indicate dependence of the implied constants.

3. The Shield Reduction

The shield reduction turns terminal game positions into a weighted lower-half antichain problem.

Theorem 3.1 (Shield reduction)**.** *Let $A \subseteq V_n$ be a maximal primitive set and let $P \subseteq A \cap U_n$. Then*

$$|A| \geq |U_n| - \beta_n(P).$$

Figure 1 illustrates the setup.

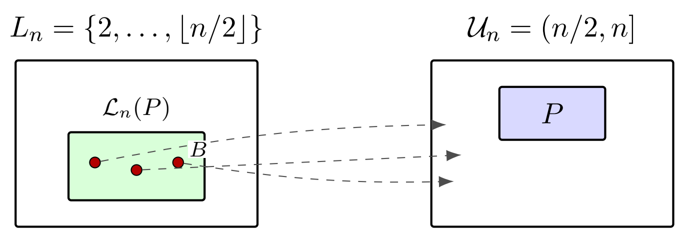

Figure 1. The shield reduction framework. A prefix $P \subseteq \mathcal{U}_n$ of Shortener-captured upper-half elements defines the residual $\mathcal{L}_n(P) = \{x \in L_n : x \nmid u \text{ for every } u \in P\}$. The weighted antichain $B \subseteq \mathcal{L}_n(P)$ maximizing $\sum_{x\in B} w_n(x)$ realizes the quantity $\beta_n(P)$. Theorem 3.1 gives $|A| \geq |\mathcal{U}_n| - \beta_n(P)$ for any terminal antichain A.

Proof. Put

$$B := A \cap L_n.$$

Then B is a divisibility antichain. Since $P \subseteq A$ and A is primitive, no $x \in B$ divides any element of P. Hence

$$B \subseteq \mathcal{L}_n(P).$$

We next determine $A \cap U_n$. We claim that

$$A \cap U_n = U_n \setminus \bigcup_{x\in B} M_n(x). \tag{3.1}$$

The inclusion "$\subseteq$" follows from primitiveness: if $u \in A \cap U_n$ and $x \in B$, then $x \nmid u$.

For the reverse inclusion, let $u \in U_n \setminus \bigcup_{x\in B} M_n(x)$. Suppose $u \notin A$. Since A is maximal, some $a \in A$ is comparable with u. If $u \mid a$, then $a \in U_n$ and $a < 2u$, so $a = u$, contradicting $u \notin A$. If $a \mid u$ and $a \in U_n$, again $a = u$. Thus $a \in L_n$, so $a \in B$, and $u \in M_n(a)$, contrary to the choice of u. This proves (3.1).

Using (3.1) and the union bound,

$$\begin{aligned}|A| &= |B| + |A \cap U_n| \\ &\geq |B| + |U_n| - \sum_{x\in B} |M_n(x)| \\ &= |U_n| - \sum_{x\in B} (|M_n(x)| - 1) \\ &= |U_n| - \sum_{x\in B} w_n(x).\end{aligned}$$

Since $B \subseteq \mathcal{L}_n(P)$ is an antichain, the final sum is at most $\beta_n(P)$. Therefore

$$|A| \geq |U_n| - \beta_n(P).$$

□

Corollary 3.2 (Shield-prefix obstruction)**.** *Fix $c > 0$. If a shield-prefix argument using Theorem 3.1 forces*

$$|A| \geq |U_n| - cn + o(n)$$

through a set $P \subseteq U_n$, then necessarily

$$|P| \geq n^{e^{-2c} - o(1)}.$$

Proof. Such an argument requires $\beta_n(P) \leq cn + o(n)$. Applying Theorem 4.1 with $|P| \leq n^\alpha$ gives

$$\frac{1}{2} \log \frac{1}{\alpha} \leq c + o(1),$$

and hence $\alpha \geq e^{-2c} - o(1)$. □

4. Lower Bounds and Conditional Slot Construction

This section proves the polynomial shield-weight obstruction and an unconditional fan-capture lower bound for the game, and then states the conditional T2 lower bound. The conditional result is deferred to Appendix A.

4.1. The polynomial shield-weight lower bound.

Theorem 4.1 (Theorem A: Polynomial shield-weight lower bound)**.** *Fix $0 < \alpha < 1$. Uniformly over all $P \subseteq U_n$ with $|P| \leq n^\alpha$,*

$$\beta_n(P) \geq \left(\frac{1}{2} \log \frac{1}{\alpha} - o(1)\right) n.$$

Proof. Fix δ with $\alpha < \delta < 1$. Let

$$Q_\delta(P) := \{p \leq n^\delta : \ p \text{ prime and } p \nmid u \text{ for every } u \in P\}.$$

Here and below, we use the Prime Number Theorem [HW08, Theorem 6] and its fixed-ratio interval form. For all sufficiently large n, every $p \leq n^\delta$ lies in L_n. The set $Q_\delta(P)$ is an antichain and is contained in $\mathcal{L}_n(P)$, so

$$\beta_n(P) \geq \sum_{p\in Q_\delta(P)} w_n(p). \tag{4.1}$$

For $p \le n^\delta$,

$$w_n(p) = \left\lfloor \frac{n}{p} \right\rfloor - \left\lfloor \frac{n}{2p} \right\rfloor - 1 = \frac{n}{2p} + O(1).$$

Since $\pi(n^\delta) = o(n)$, (4.1) gives

$$\beta_n(P) \ge \frac{n}{2} \sum_{p \in Q_\delta(P)} \frac{1}{p} - o(n). \tag{4.2}$$

Let

$$C_\delta(P) := \{p \le n^\delta : \ p \text{ prime and } p \mid u \text{ for some } u \in P\}.$$

Then $Q_\delta(P)$ and $C_\delta(P)$ partition the primes up to n^δ.

We first bound the reciprocal mass of $C_\delta(P)$. For each $p \in C_\delta(P)$ choose one $u(p) \in P$ divisible by p. For any fixed $u \in P$, the primes assigned to u are distinct divisors of u, so

$$\sum_{p\colon u(p)=u} \log p \le \log u \le \log n.$$

Summing over $u \in P$ gives

$$\sum_{p \in C_\delta(P)} \log p \le |P| \log n \le n^\alpha \log n. \tag{4.3}$$

Among all finite prime sets satisfying a log-budget $\sum_{p\in S} \log p \le B$, the reciprocal sum $\sum_{p \in S} 1/p$ is maximized by taking the smallest primes first. Indeed, if S contains a prime q and omits a smaller prime $p < q$, then replacing q by p decreases the log-cost and increases the reciprocal sum. Iterating this exchange proves the claim.

Let z_n be the largest value for which

$$\vartheta(z_n) = \sum_{p \le z_n} \log p \le n^\alpha \log n.$$

By (4.3) and the exchange argument,

$$\sum_{p \in C_\delta(P)} \frac{1}{p} \le \sum_{p \le z_n} \frac{1}{p}. \tag{4.4}$$

Chebyshev's theorem $\vartheta(x) \sim x$ [HW08, Theorem 7] gives

$$z_n = (1 + o(1)) n^\alpha \log n,$$

and hence

$$\log\log z_n = \log\log n + \log \alpha + o(1).$$

Mertens' theorem for primes [HW08] yields

$$\begin{aligned}
\sum_{p \in Q_\delta(P)} \frac{1}{p} &= \sum_{p \le n^\delta} \frac{1}{p} - \sum_{p \in C_\delta(P)} \frac{1}{p} \\
&\ge \sum_{p \le n^\delta} \frac{1}{p} - \sum_{p \le z_n} \frac{1}{p} \\
&= \log\log(n^\delta) - \log\log z_n + o(1) \\
&= \log \frac{\delta}{\alpha} + o(1).
\end{aligned}$$

Substituting this into (4.2) gives

$$\beta_n(P) \ge \left(\frac{1}{2} \log \frac{\delta}{\alpha} - o(1) \right) n.$$

Letting $\delta \uparrow 1$ proves the theorem. ☐

Corollary 4.2 (Barrier exponent)**.** *For every fixed $c > 0$,*

$$k_c^\star(n) := \min\{|P| : \ P \subseteq U_n,\ \beta_n(P) \le cn\} \ge n^{e^{-2c}-o(1)}.$$

Proof. If $|P| \le n^\alpha$ and $\beta_n(P) \le cn$, then Theorem 4.1 gives

$$\frac{1}{2}\log\frac{1}{\alpha} \le c + o(1).$$

Thus $\alpha \ge e^{-2c} - o(1)$, which is the claimed exponent lower bound. ☐

4.2. **An unconditional fan-capture lower bound.** The trivial unconditional lower bound for the game is $L(n) \ge (1+o(1))n/\log n$: in any maximal primitive subset $A \subseteq \{2, \dots, n\}$, maximality forces A to contain, for each prime $p \in (\sqrt{n}, n]$, some multiple of p; and two distinct primes in $(\sqrt{n}, n]$ have no common multiple in $\{2, \dots, n\}$, giving $|A| \ge \pi(n) - \pi(\sqrt{n}) \sim n/\log n$ [Erd26]. The theorem below improves this baseline by a factor of order $\log\log n$, giving the first unconditional lower bound of strictly larger order.

The shield obstruction above does not itself give a game lower bound. We now record a direct unconditional lower-bound strategy for Prolonger. The mechanism is different: Prolonger first makes a large reciprocal-mass set of small primes unavailable, then uses the resulting fan graph of upper-half products.

Lemma 4.3 (Maximum-degree right capture)**.** *Let $G = (L \sqcup R, E)$ be a finite bipartite graph. Consider the following finite capture game on the live graph. Maker moves first. On each Maker move, Maker chooses a right vertex of maximum positive live degree, captures all currently live edges incident to it, and removes that right vertex and its incident live edges. On each Breaker move, Breaker may delete one right vertex and all live edges incident to it, delete one live edge, or do nothing. The game stops when no live edge remains.*

Let C be the number of edges captured by Maker, and let X be the number of edges individually deleted by Breaker. Then, regardless of Breaker's play,

$$C + X \ge \frac{|E|}{2}.$$

Proof. Index Maker's nontrivial moves by i. Let d_i be the live degree of the right vertex captured by Maker on move i. Thus this move contributes d_i to C. Since Maker chose a maximum-degree right vertex, after Maker removes it every remaining right vertex has live degree at most d_i. Therefore, if Breaker replies by deleting a right vertex in the same round, that deletion removes at most d_i live edges. If Breaker instead deletes one edge or does nothing, the number of right-deleted edges in that round is 0.

Let D_{right} be the total number of edges deleted through right-vertex deletions. Summing the preceding round-by-round bound gives

$$D_{\text{right}} \le \sum_i d_i = C.$$

Every original edge is eventually in exactly one of three classes: captured by Maker, deleted as part of a right-vertex deletion, or individually deleted by Breaker. Hence

$$|E| = C + D_{\text{right}} + X \le 2C + X \le 2(C + X),$$

which proves the claim. ☐

Theorem 4.4 (Unconditional fan-capture lower bound)**.** *For every fixed $0 < \delta < 1/2$ and every $\varepsilon > 0$, for all sufficiently large n, Prolonger has a strategy forcing*

$$L(n) \ge \left(\frac{1}{8} - \varepsilon\right)\frac{n\log\log n}{\log n}.$$

Equivalently,

$$L(n) \geq \left(\frac{1}{8} - o(1)\right) \frac{n \log\log n}{\log n}.$$

Proof. Fix $0 < \delta < 1/2$ and put $y := n^{\delta}$. Let

$$\mathcal{A} := \{a \leq y : \ a \in \mathbb{P}, \ a \text{ odd}\}.$$

For each $a \in \mathcal{A}$, define the right-prime interval

$$J_a := \left(\frac{n}{2a}, \frac{n}{a}\right] \cap \mathbb{P}.$$

If $b \in J_a$, then $ab \in U_n$. Also $b > y$ for all sufficiently large n, because $b > n/(2y) = \frac{1}{2}n^{1-\delta}$ and $\delta < 1/2$.

We shall use the fan graph with left vertices in $\mathcal{A}$, right vertices the primes appearing in the intervals J_a, and edges (a, b) representing the upper-half target ab.

Phase 1: activating small primes. As long as some prime in $\mathcal{A}$ is still a legal singleton at the start of a Prolonger turn, Prolonger takes the smallest such prime a and plays a legal target ab with $b \in J_a$. Such a target exists. Indeed, before any Phase 1 Prolonger move there have been $O(|\mathcal{A}|)$ previous Phase 1 moves, while uniformly for $a \in \mathcal{A}$ the prime number theorem gives

$$|J_a| = (1 + o_{\delta}(1)) \frac{n}{2a \log(n/a)} \gg_{\delta} \frac{n^{1-\delta}}{\log n}.$$

This is larger than $|\mathcal{A}| = O_{\delta}(n^{\delta}/\log n)$ because $\delta < 1/2$. If a is still legal, a previously chosen number can make a candidate ab illegal only by being b itself or ab itself; each previous move rules out at most one candidate b. Therefore some legal ab remains.

Phase 1 ends when no prime in $\mathcal{A}$ is a legal singleton. If this happens immediately after a Prolonger move, include Shortener's next move, if there is one, in the Phase 1 damage accounting below before Phase 2 starts. This convention only adds one Shortener move.

Define

$$A_0 := \{a \in \mathcal{A} : \text{ during Phase 1, } a \text{ is made illegal by a chosen proper multiple}\}$$

and

$$\mathcal{D} := \{a \in \mathcal{A} : \text{ Shortener chooses the singleton } a \text{ during Phase 1}\}.$$

Prolonger never plays a singleton from $\mathcal{A}$, and a prime singleton can become illegal only by being chosen itself or by having a proper multiple chosen. Thus

$$\mathcal{A} = A_0 \sqcup \mathcal{D}.$$

Let $P \subseteq A_0$ be the set of primes activated by Prolonger's Phase 1 moves. Pair each $a' \in \mathcal{D}$ with the immediately preceding Prolonger activation $a \in P$. The pairs are distinct because turns alternate. Since Prolonger always activates the smallest currently live prime, and since legality only decreases as play proceeds, $a \leq a'$. Therefore

$$\sum_{a' \in \mathcal{D}} \frac{1}{a'} \leq \sum_{a \in P} \frac{1}{a} \leq \sum_{a \in A_0} \frac{1}{a}.$$

Using $\mathcal{A} = A_0 \sqcup \mathcal{D}$ and Mertens' theorem for primes,

$$\sum_{a \in A_0} \frac{1}{a} \geq \frac{1}{2} \sum_{a \in \mathcal{A}} \frac{1}{a} = \left(\frac{1}{2} - o(1)\right) \log\log n. \tag{4.5}$$

Phase 1 damage and the surviving fan. Let $E^{\star}$ be the raw fan-edge set over A_0:

$$E^{\star} := \{(a, b) : a \in A_0, \ b \in J_a\}.$$

The same uniform prime number theorem estimate gives

$$
\begin{aligned}
|E^\star| &= \sum_{a\in A_0} |J_a| \\
&\geq \left(\frac{1}{2} - o_\delta(1)\right) \frac{n}{\log n} \sum_{a\in A_0} \frac{1}{a} \\
&\geq \left(\frac{1}{4} - o_\delta(1)\right) \frac{n\log\log n}{\log n},
\end{aligned}
\tag{4.6}
$$

by (4.5).

We now remove from $E^\star$ the edges whose targets have already become unavailable before the first Phase 2 Prolonger move, and call the remaining edge set E_0. If $a \in A_0$, the singleton a was not chosen. Since $ab > n/2$, the target ab has no proper multiple in $\{2, \ldots, n\}$, and its only relevant proper divisors are a and b. Hence a Phase 1 move can destroy a raw fan edge (a, b) only by choosing the right singleton b, which removes a whole right star, or by choosing the exact target ab, which removes one edge.

There are $O(|\mathcal{A}|)$ Phase 1 moves, including the possible intervening Shortener move. Right-singleton choices therefore destroy at most $O(|\mathcal{A}|^2)$ raw fan edges, and exact-target choices destroy at most $O(|\mathcal{A}|)$ raw fan edges. Since

$$|\mathcal{A}| = O_\delta\left(\frac{n^\delta}{\log n}\right),$$

the total damage is

$$O_\delta\left(\frac{n^{2\delta}}{\log^2 n}\right) = o_\delta\left(\frac{n}{\log n}\right) = o_\delta\left(\frac{n\log\log n}{\log n}\right).$$

Consequently

$$|E_0| \geq \left(\frac{1}{4} - o_\delta(1)\right) \frac{n\log\log n}{\log n}. \tag{4.7}$$

Phase 2: capturing right stars. On the surviving fan graph $G_0 = (A_0 \sqcup R_0, E_0)$, Prolonger follows the maximum-degree right-capture strategy from Lemma 4.3. That is, on each Phase 2 turn he chooses a right prime b of maximum positive current degree and plays one legal target $a_0 b$ incident to it.

This move is legal. The left singleton a_0 is already illegal from Phase 1, but it was not chosen; the right singleton b has not been chosen, because otherwise the whole right star would have been removed from E_0; and the target $a_0 b$ itself is still available by the definition of E_0. Since $a_0 b > n/2$, no proper multiple of it is on the board, and the only relevant proper divisors are a_0 and b.

Shortener's reply has one of the three effects allowed in Lemma 4.3: choosing a right singleton deletes a right vertex and its live star; choosing an exact fan target ab deletes that one live edge if it is still live, and is a do-nothing move relative to the live graph if ab was already captured in an earlier right-star capture by Prolonger; any other legal move deletes no live fan edge. Therefore the lemma applies. Let C be the number of fan edges captured by Prolonger's right-star moves in Phase 2, and let X be the number of live fan edges individually deleted by Shortener as exact targets in Phase 2. Then

$$C + X \geq \frac{|E_0|}{2}. \tag{4.8}$$

It remains to translate captured graph edges back to actual game moves. When Prolonger captures a right prime b by playing some $a_0 b$, the capture-move edge (a_0, b) is realized as Prolonger's current game move. For every *other* edge (a, b) live at that capture time, with $a \neq a_0$, the singleton a is already illegal from Phase 1 but was not chosen, and the target ab has not yet been chosen.

Moreover ab has no proper multiple in $\{2, \dots, n\}$ and its only relevant proper divisors are a and b; any proper multiple of a that was chosen in Phase 1 is not comparable with this surviving target. Therefore ab remains legal until ab itself is chosen; terminal maximality forces it to be chosen eventually. Distinct fan targets all lie in U_n, so they are mutually incomparable unless they are equal, and the graph edges are distinct. Thus the captured edges account for C distinct eventual game moves (one of which, $a_0 b$, is Prolonger's own capture play), while the individually deleted edges counted by X are themselves distinct game moves realized by Shortener's exact-target choices. Hence

$$L(n) \geq C + X \geq \frac{|E_0|}{2} \geq \left(\frac{1}{8} - o_\delta(1)\right) \frac{n \log\log n}{\log n},$$

by (4.7) and (4.8). Since δ was fixed, this proves the theorem. □

4.3. **The conditional T2 theorem.** The preceding fan-capture theorem gives an unconditional game lower bound. We now record a different, potentially larger lower-bound mechanism as a *conditional target theorem*: the unconditional version remains open, and the hypothesis below is stated explicitly so that any future progress can be measured against a precise finite-combinatorial target rather than a folklore belief. The mechanism uses residual upper-half targets of the form acb, where a, c are small primes and b lies in a corresponding large-prime interval. The resulting residual game is modeled by a scored rank-three slot hypergraph.

The finite potential argument which one would like to use in this slot hypergraph is not true for arbitrary finite capture games. The counterexample in Proposition A.1 shows that a unique max-gain Maker move need not dominate every subsequent Breaker deletion, and Proposition A.2 shows that the general safe-edge property already fails inside the residual arithmetic games. We therefore state the lower bound conditionally on a restricted property for the states generated by the specific activation and residual strategies.

Definition 4.5 (Restricted safe-edge hypothesis)**.** In each finite graph or scored rank-three hypergraph state actually reached by the T2 activation strategy and subsequent residual slot construction, assume that whenever a positive-weight live edge exists, Maker has a legal live edge f such that, after Maker plays f, every legal Breaker reply leaves the scaled potential Q, as defined in the finite slot-game setup of Appendix A, at least as large as it was before Maker's move.

Remark 4.6. The narrowing to strategy-generated states is specifically engineered to avoid the counterexample of Proposition A.2: the restricted hypothesis is therefore the narrowest form compatible with the paper's own refutation of the general safe-edge property, and its plausibility is an open target.

Theorem 4.7 (Conditional T2 lower bound)**.** *Assume the restricted safe-edge hypothesis for the finite graph and residual slot-hypergraph states generated by the construction in Appendix A. Then for every fixed* $0 < \delta < 1/4$ *there is a constant* $c_\delta > 0$ *such that*

$$L(n) \geq c_\delta \frac{n(\log\log n)^2}{\log n}$$

for all sufficiently large n*.*

Proof. See Proposition A.9 in Appendix A. The full proof assembles Propositions A.5 to A.8 under the restricted safe-edge hypothesis of Definition 4.5. □

5. The Exact 5/24 First-Hit Cover

This section determines the minimum number of upper-half elements needed to hit every lower-half element by divisibility.

Theorem 5.1 (First-hit cover constant)**.** *Let*

$$\tau(n) := \min\{|C| : \ C \subseteq U_n, \ \forall x \in L_n \ \exists u \in C \ \textit{with } x \mid u\}.$$

Then

$$\tau(n) = \frac{5}{24}n + O(1).$$

More precisely, the cover

$$H_n = \{u \in U_n : \ u \equiv 2 \pmod 4\} \cup \{u \in U_n : \ u > 2n/3, \ u \equiv 0 \pmod 4\}$$

has size $5n/24 + O(1)$*, and there is a matching packing of lower-half witnesses of the same size up to* $O(1)$*.*

See Figure 2 for the explicit construction.

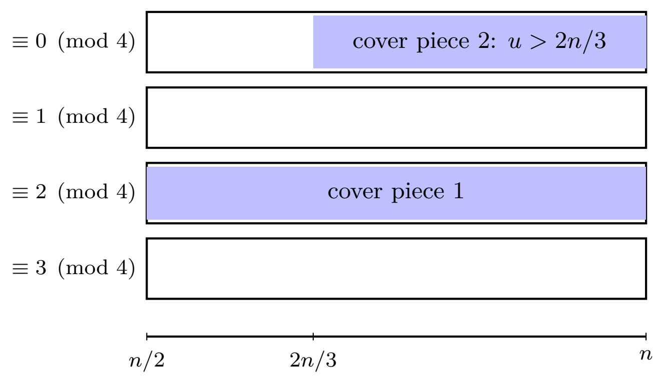

$$|H_n| = \tfrac{5n}{24} + O(1)$$

Figure 2. The exact first-hit cover H_n from Theorem 1.6. Upper-half elements are stratified by residue mod 4. The cover consists of all $u \equiv 2 \pmod 4$ (entire band) together with $u \equiv 0 \pmod 4$ restricted to $u > 2n/3$ (right portion of the $\equiv 0$ band). Residue classes $\equiv 1$ and $\equiv 3$ are odd and require no cover element beyond their own presence. The matching lower-half packing (not shown) gives $\tau(n) = 5n/24 + O(1)$.

Proof. We prove matching upper and lower bounds.

For the upper bound, define

$$H_n := \{u \in U_n : \ u \equiv 2 \pmod 4\} \cup \{u \in U_n : \ u > 2n/3, \ u \equiv 0 \pmod 4\}.$$

We show that every $x \in L_n$ divides some element of H_n.

If $x > n/3$, then $2x \in U_n$. If x is odd, then $2x \equiv 2 \pmod 4$; if x is even, then $2x \equiv 0 \pmod 4$ and $2x > 2n/3$. Hence $2x \in H_n$.

If $n/4 < x \le n/3$, then for odd x the element $2x$ lies in U_n and is 2 modulo 4. For even x, the element $3x$ satisfies $3x > 3n/4 > 2n/3$ and $3x \le n$; it is even, and is therefore either 2 or 0 modulo 4, with the latter case covered by the condition $3x > 2n/3$. Thus $3x \in H_n$.

If $n/6 < x \le n/4$, then $4x \in U_n$, $4x > 2n/3$, and $4x \equiv 0 \pmod 4$, so $4x \in H_n$.

Finally suppose $x \le n/6$. The interval

$$\left(\frac{2n}{3x}, \frac{n}{x}\right]$$

has length $n/(3x) \ge 2$, and hence contains an even integer k. Then $u = kx$ satisfies $2n/3 < u \le n$ and is even. If $u \equiv 2 \pmod 4$ then $u \in H_n$ by the first clause; if $u \equiv 0 \pmod 4$ then $u \in H_n$ by the second clause. In either case $x \mid u$.

The two clauses defining H_n are disjoint. Therefore

$$\begin{aligned} |H_n| &= \#\{u \in (n/2, n] : u \equiv 2 \pmod 4\} \\ &\quad + \#\{u \in (2n/3, n] : u \equiv 0 \pmod 4\} \\ &= \frac{n}{8} + \frac{n}{12} + O(1) = \frac{5n}{24} + O(1). \end{aligned}$$

Thus $\tau(n) \le 5n/24 + O(1)$.

For the lower bound, define

$$P_n := \{m : \ n/3 < m \le n/2\} \cup \{m : \ n/4 < m \le n/3, \ m \text{ odd}\}.$$

Then

$$|P_n| = \frac{n}{6} + \frac{1}{2}\left(\frac{n}{3} - \frac{n}{4}\right) + O(1) = \frac{5n}{24} + O(1).$$

The set P_n is an antichain. Indeed, if $x < y$ and $x \mid y$, then $y \ge 2x$. If $y \le n/3$, then $x \le n/6 < n/4$, impossible for $x \in P_n$. If $n/3 < y \le n/2$, then $x \le n/4$, again impossible for $x \in P_n$.

We claim that any $u \in U_n$ is divisible by at most one element of P_n. If $x \in P_n$ divides u, then $x > n/4$ and $u \le n$, so

$$1 < \frac{u}{x} < 4.$$

Thus u/x is either 2 or 3. If two distinct elements of P_n divide u, they must be $u/2$ and $u/3$. In particular $6 \mid u$. But then $u/3$ is even. It cannot lie in the first component $(n/3, n/2]$ because $u/3 \le n/3$, and it cannot lie in the second component because that component requires oddness. Hence at most one element of P_n divides u.

Now let $C \subseteq U_n$ be any cover of L_n. Each $x \in P_n$ must divide some element of C, and the preceding paragraph shows that different $x \in P_n$ require different elements of C. Therefore

$$|C| \ge |P_n| = \frac{5n}{24} + O(1).$$

This proves the matching lower bound. □

6. Intermediate Upper Bounds

The main upper bound in Section 7 refines a simpler odd-prime-prefix method. We present the two intermediate versions because their proofs isolate the compression and Bonferroni ideas used later.

6.1. Odd-part compression.

The following elementary lemma is used throughout the upper-bound arguments.

Lemma 6.1 (Odd-part injection)**.** *Let A be a divisibility antichain. The map*

$$\varphi(x) := \frac{x}{2^{v_2(x)}}$$

is injective on A.

Proof. If $\varphi(x) = \varphi(y) = m$, then $x = 2^a m$ and $y = 2^b m$. If $a < b$, then $x \mid y$; if $b < a$, then $y \mid x$. Since A is an antichain, both cases force $x = y$. □

6.2. **The 13/36 bound.**

Theorem 6.2. *Shortener has a strategy forcing*

$$L(n) \le \left(\frac{13}{36} + o(1)\right) n.$$

Proof. Let

$$k := \left\lfloor \frac{\sqrt{n}}{\log n} \right\rfloor .$$

For her first k turns, Shortener plays the smallest legal odd prime; after that she plays arbitrarily. Let

$$q_1 < q_2 < \cdots < q_k$$

be the odd primes she obtains.

We first bound these primes. Fix $\varepsilon > 0$. If, before Shortener's jth prefix move, every odd prime up to

$$X = (3/2 + \varepsilon) j \log n$$

were blocked, then the total odd-prime Chebyshev mass up to X would be covered by the earlier Shortener primes together with the odd prime divisors of Prolonger's first j moves. Each Prolonger move contributes at most $\log n$ of prime-divisor log-mass. By induction on j, the previous Shortener primes contribute at most

$$\sum_{i<j} \log q_i \le \left(\frac{1}{2} + o(1)\right) j \log n,$$

because $j \le \sqrt{n}/\log n$. Thus the blocked log-mass is at most $(3/2 + o(1)) j \log n$, while the prime number theorem gives

$$\vartheta_{\text{odd}}(X) = (3/2 + \varepsilon + o(1)) j \log n,$$

a contradiction for large n. Hence

$$q_j \le (3/2 + \varepsilon) j \log n$$

uniformly for $j \le k$. For this fixed ε,

$$\sum_{j=1}^{k} \frac{1}{q_j} \ge \frac{1}{3 + 2\varepsilon} - o_\varepsilon(1).$$

Since the strategy and the preceding estimate hold for every $\varepsilon > 0$, we may let $\varepsilon \downarrow 0$ after taking $n \to \infty$, and therefore

$$\sum_{j=1}^{k} \frac{1}{q_j} \ge \frac{1}{3} - o(1). \tag{6.1}$$

Let A be the final antichain. All elements other than the prefix primes have odd part not divisible by any q_j. By Lemma 6.1,

$$|A| \le k + N_D(n),$$

where $D = \{q_1, \ldots, q_k\}$ and $N_D(n)$ is the number of odd integers $m \le n$ not divisible by any prime in D.

Choose t minimal such that

$$s_t := \sum_{j \le t} \frac{1}{q_j} \ge \frac{1}{3} - o(1),$$

and put $E = \{q_1, \ldots, q_t\}$. Since every q_j is an odd prime, $1/q_j \le 1/3$, so

$$s_t \in \left[\frac{1}{3} - o(1), \frac{2}{3} + o(1)\right] \subset [0, 1]$$

for large n. Because $E \subseteq D$, we have $N_D(n) \le N_E(n)$.

Second-order Bonferroni on the events $q \mid m$ over odd $m \le n$ gives

$$N_E(n) \le \frac{n}{2}\left(1 - s_t + \frac{s_t^2}{2}\right) + O(t^2).$$

Here $t \le k$, so $O(t^2) = O(n/\log^2 n) = o(n)$. The function $f(s) = 1 - s + s^2/2$ is decreasing on $[0, 1]$, and therefore

$$f(s_t) \le f(1/3 - o(1)) = \frac{13}{18} + o(1).$$

Consequently

$$|A| \le k + N_D(n) \le \frac{13}{36}n + o(n).$$

□

6.3. **The $5/16$ bound.**

Theorem 6.3. *Shortener has a strategy forcing*

$$L(n) \le \left(\frac{5}{16} + o(1)\right) n.$$

Proof. Fix $A_\star > 2$ and set

$$k := \left\lfloor \frac{n}{2A_\star \log n} \right\rfloor.$$

Shortener plays the smallest legal odd prime for her first k turns. (Legality of such a prefix for all sufficiently large n is proved in Lemma 7.1; the argument there depends only on board size, not on the odd-part compression used in this section.) The same Chebyshev log-mass induction, now with $j \le n/(2A_\star \log n)$, gives

$$q_j \le A_\star j \log n \qquad (1 \le j \le k).$$

Indeed, if all odd primes below $A_\star j \log n$ were blocked, Prolonger's first j moves would contribute at most $j \log n$, and the previous Shortener primes would contribute at most $(1 + o(1))j \log n$; this is less than $A_\star j \log n$ because $A_\star > 2$.

It follows that

$$\sum_{j=1}^{k} \frac{1}{q_j} \ge \frac{1}{A_\star} - o(1).$$

As in the proof of Theorem 6.2, use the odd-part injection and truncate the prime set at the first partial sum

$$s_t \ge \frac{1}{A_\star} - o(1).$$

Since $1/q_j \le 1/3$ and $A_\star > 2$, the partial sum remains in $[0, 1]$. The second-order Bonferroni estimate gives

$$L(n) \le \frac{n}{2}\left(1 - \frac{1}{A_\star} + \frac{1}{2A_\star^2}\right) + o(n).$$

For every fixed $A_\star > 2$ the strategy above is valid. Letting $A_\star \downarrow 2$ after taking $n \to \infty$ yields

$$L(n) \le \frac{n}{2}\left(1 - \frac{1}{2} + \frac{1}{8}\right) + o(n) = \frac{5}{16}n + o(n).$$

The pair-intersection error is $o(n)$ because only prime pairs with product at most n contribute an $O(1)$ floor error, and

$$\#\{p < q :\ pq \le n\} \ll \sum_{p \le \sqrt{n}} \pi(n/p) \ll \frac{n}{\log n} \sum_{p \le \sqrt{n}} \frac{1}{p} = o(n).$$

□

7. The Main Upper Bound

In this section we prove the paper's headline upper bound

$$L(n) \le \left(\frac{\mathcal{W}_4}{2} + o(1)\right) n < 0.19n.$$

The proof uses the following explicit strategy. Shortener plays the smallest legal odd prime for a long initial prefix of

$$K = \left\lfloor \frac{(1-\varepsilon)n}{2\log n} \right\rfloor$$

turns. This prefix is available for all sufficiently large n.

Lemma 7.1 (Existence of the odd-prime prefix)**.** *Fix $0 < \varepsilon < 1$ and*

$$K = \left\lfloor \frac{(1-\varepsilon)n}{2\log n} \right\rfloor.$$

For all sufficiently large n, before each of Shortener's first K turns there is a legal odd prime. Hence the smallest-legal-odd-prime strategy can be followed for a prefix of K Shortener moves.

Proof. Since the board is $\{2, 3, \dots, n\}$, no legal move equals 1. Consider the moment before Shortener's jth move in this prefix, with $1 \le j \le K$. At most $2j - 1$ previous moves have occurred. A previous move $x \le n$ can make at most one odd prime $p > \sqrt{n}$ illegal: if two such primes divided x, then $x > n$, while the only other way for x to be comparable with such a prime is $x = p$ itself. Thus at most $2j - 1$ odd primes in $(\sqrt{n}, n]$ have been made illegal.

By the prime number theorem,

$$\pi(n) - \pi(\sqrt{n}) \sim \frac{n}{\log n} > 2K$$

for all sufficiently large n. Since $2j - 1 < 2K$, at least one odd prime $p \in (\sqrt{n}, n]$ remains legal before Shortener's jth prefix move. In particular the game has not ended, and Shortener's rule has a legal odd prime to play. ☐

Let

$$q_1 < q_2 < \cdots < q_K$$

be the odd primes she plays during that prefix. The argument has four layers:

(1) a one-sided lower-profile estimate for the counting function of the q_j;
(2) a monotone envelope and inversion producing a comparison sequence $b_j \ge q_j$ with explicit moment limits;
(3) a fourth-order Bonferroni comparison theorem for arbitrary increasing odd-prime comparison sequences;
(4) a prime-rounding bridge that replaces the real sequence b_j by actual odd primes p_j, using the key half-density slack supplied by the prime-rounding bridge of Theorem 7.20.

The assembly from Lemma 7.4 through Lemma 7.17 is the longest prose stretch in the paper and the part most sensitive to auditing. Its arithmetic kernels—envelope properties and flat-block mass estimates—have zero-sorry Lean artifacts in the files `Envelope.lean` and `FlatMass.lean` of the `Round15Bonferroni4` project; the assembly itself is prose, with the moment-convergence hypotheses feeding the zero-sorry endgame reduction in `Target.lean`.

7.1. **A lower-profile envelope for Shortener's primes.** Write

$$S_K(X) := \#\{j \le K : q_j \le X\}, \qquad u := \frac{\log X}{\log n}.$$

We establish a one-sided lower envelope for the counting function of Shortener's captured primes. On the interval

$$u \in \left(\frac{1}{h+1}, \frac{1}{h}\right],$$

the lower-profile function is

$$\rho(u) = \frac{1}{(\lfloor 1/u \rfloor + 1)u} = \frac{1}{(h+1)u}.$$

Figure 3 plots this piecewise-hyperbolic profile on $(0, 1]$.

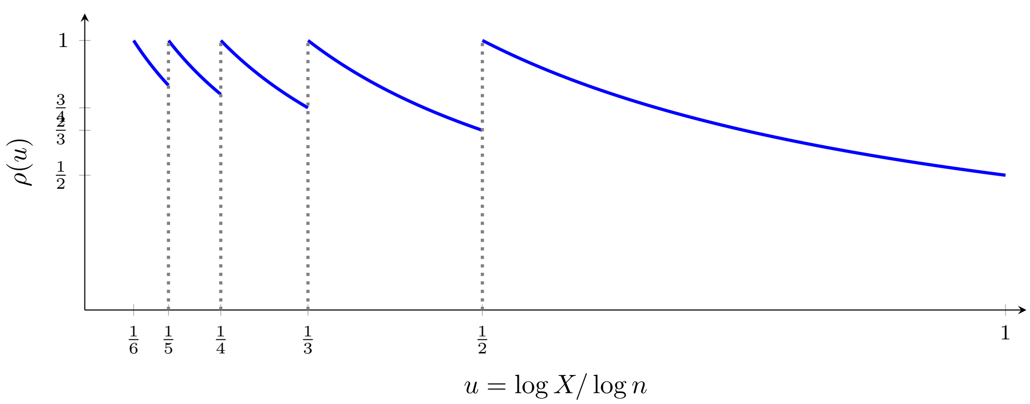

FIGURE 3. The lower-profile function $\rho(u) = 1/((\lfloor 1/u \rfloor + 1)u)$ on $(0, 1]$. On each interval $(1/(h + 1), 1/h]$ with $h \ge 1$, ρ is the strictly decreasing hyperbola $u \mapsto 1/((h + 1)u)$, jumping up to 1 at each left endpoint $u = 1/(h + 1)^+$ and descending to $h/(h + 1)$ at $u = 1/h$. The monotone right-continuous cumulative envelope built in Lemma 7.4 and inverted in Lemma 7.5 is derived from this profile.

Lemma 7.2 (Local prime-count per range)**.** *Let Shortener follow the smallest-legal-odd-prime strategy. Fix an integer $h \ge 1$ and real numbers $2 \le Y < X \le n$ with*

$$Y^{h+1} > n.$$

Assume the odd-prime strategy is available until the first time her next odd-prime move would exceed X, and let $S(X)$ be the number of odd primes at most X which Shortener has played by that stopping time. Then

$$S(X) \ge \frac{\pi(X) - \pi(Y)}{h+1} - O_h(1).$$

In particular, if $n^{1/(h+1)} < Y < X \le n^{1/h}$, then every prime $p \in (Y, X]$ lies in the packet regime $p^h \le n < p^{h+1}$ and the same lower bound holds.

Proof. Shortener's odd-prime moves are increasing: once a prime has become illegal it cannot become legal later, and the strategy always chooses the smallest legal odd prime. At the stopping time in the definition of $S(X)$, no odd prime $p \le X$ is still legal. Therefore every odd prime $p \in (Y, X]$ has one of two forms: either Shortener herself has played p, or some earlier Prolonger move was divisible by p.

Let

$$M := \#\{p \in \mathbb{P} : \ Y < p \le X, \ p \text{ odd}\}.$$

The number of primes counted by M that Shortener has played is at most $S(X)$. There are at most $S(X)+1$ Prolonger moves before the stopping time. Each such move is an integer at most n, so it is divisible by at most h distinct primes from $(Y,X]$: otherwise it would be at least $Y^{h+1}>n$. Hence

$$M \le S(X)+h(S(X)+1)=(h+1)S(X)+h.$$

Thus

$$S(X) \ge \frac{M-h}{h+1}.$$

Finally $M=\pi(X)-\pi(Y)+O(1)$, the $O(1)$ accounting only for the prime 2 and endpoint conventions. This gives the claimed bound. If $n^{1/(h+1)}<Y<X\le n^{1/h}$ and $p\in(Y,X]$, then $p^h \le X^h \le n$ while $p^{h+1}>Y^{h+1}>n$, so this is exactly the stated packet regime. □

Proposition 7.3 (Local density away from breakpoints)**.** *Fix $H\ge 2$, and set*

$$\tau_H := \frac{1}{8H^2}.$$

For $1\le h\le H-1$, define

$$\alpha_h := \frac{1}{h+1}+\tau_H, \qquad \beta_h := \frac{1}{h}-\tau_H, \qquad I_h := [\alpha_h,\beta_h].$$

Then there exists a function $\xi_H(n)\to 0$ such that, for all sufficiently large n, for every $1\le h\le H-1$ and every $u\in I_h$,

$$S_K(n^u) \ge \big(1-\xi_H(n)\big)\frac{n^u}{(h+1)u\log n}.$$

Proof. The intervals I_h are nonempty, because

$$\beta_h-\alpha_h = \frac{1}{h}-\frac{1}{h+1}-2\tau_H \ge \frac{1}{H(H-1)}-\frac{1}{4H^2}>0.$$

Fix $1\le h\le H-1$ and $u\in I_h$, and put

$$\eta := \frac{\tau_H}{2}, \qquad X:=n^u, \qquad Y:=n^{u-\eta}.$$

Since $u\in I_h$,

$$u-\eta \ge \frac{1}{h+1}+\frac{\tau_H}{2} > \frac{1}{h+1}, \qquad u\le \frac{1}{h}-\tau_H<\frac{1}{h}.$$

Hence the entire interval $[u-\eta,u]$ stays inside $\left(\frac{1}{h+1},\frac{1}{h}\right)$. Thus $Y^{h+1}>n$, and Lemma 7.2 yields

$$S(X) \ge \frac{\pi(X)-\pi(Y)}{h+1}-O_H(1),$$

where $S(X)$ denotes the number of odd primes played by Shortener up to the first time her move exceeds X.

Assuming the prime prefix exists for all K Shortener moves, as guaranteed for large n by Lemma 7.1, we claim that $S(X)=S_K(X)$ for all $u\in\bigcup_{h=1}^{H-1} I_h$ once n is large enough. Indeed, $u\le\beta_1=1-\tau_H$, so by the prime number theorem,

$$\pi(n^u) \ll \frac{n^{1-\tau_H}}{\log n} = o\!\left(\frac{n}{\log n}\right)=o(K).$$

Therefore all odd primes at most $X=n^u$ are exhausted before the prefix of length K, and the local counting function $S(X)$ coincides with the truncated counting function $S_K(X)$.

Substituting $X=n^u$ and $Y=n^{u-\eta}$ gives

$$S_K(n^u) \ge \frac{\pi(n^u)-\pi(n^{u-\eta})}{h+1}-O_H(1).$$

Now u ranges over the fixed compact interval I_h, so the prime number theorem is uniform there:

$$\pi(n^u) = \big(1 + o_H(1)\big)\frac{n^u}{u \log n}.$$

Also $u - \eta \le u - \tau_H/2$, hence

$$\frac{\pi(n^{u-\eta})}{n^u/(u\log n)} \ll_H u\,n^{-\eta} = o_H(1)$$

uniformly for $u \in I_h$. Since there are only finitely many blocks $1 \le h \le H-1$, we may absorb the error terms into a single function $\xi_H(n) \to 0$, valid simultaneously for every h and every $u \in I_h$. This gives

$$S_K(n^u) \ge \big(1 - \xi_H(n)\big)\frac{n^u}{(h+1)u\log n},$$

as claimed. □

7.2. **Moment constants and the Bonferroni coefficient.** Define

$$J_r := \frac{1}{r!}\int_{\substack{u_1+\cdots+u_r\le 1\\ u_i\in(0,1]}} \prod_{i=1}^{r} \rho(u_i)\, du_1 \cdots du_r,$$

and

$$\mathcal{W}_4 := 1 - J_1 + J_2 - J_3 + J_4.$$

The interval-arithmetic certification in Proposition B.1 gives

$$\frac{\mathcal{W}_4}{2} \le 0.1897123371 < 0.19.$$

For any increasing sequence $b_1 \le \cdots \le b_K$ of positive reals, define the truncated factorial-moment sums

$$T_r^{(b)}(n) := \sum_{\substack{1\le j_1<\cdots<j_r\le K\\ b_{j_1}\cdots b_{j_r}\le n}} \frac{1}{b_{j_1}\cdots b_{j_r}} \qquad (1 \le r \le 4).$$

7.3. **Monotone envelope and inversion.** The lower-profile estimate from Proposition 7.3 is piecewise-defined, so one must build a monotone global comparison sequence before applying the Bonferroni comparison.

Lemma 7.4 (Right-continuous cumulative envelope)**.** *Fix $H \ge 2$, and define $\tau_H, \alpha_h, \beta_h, I_h$ as in Proposition 7.3. For every sufficiently large n there is a monotone right-continuous step function $C_{H,n} : [2, n] \to \{0, 1, \ldots, K\}$ such that*

$$C_{H,n}(X) \le S_K(X) \qquad (2 \le X \le n).$$

More precisely, with $\xi_H(n) \to 0$ chosen from Proposition 7.3, put

$$A_{n,h}(u) := \big(1 - \xi_H(n)\big)\frac{n^u}{(h+1)u\log n} \qquad (u \in I_h),$$

and

$$L_h := \lfloor A_{n,h}(\alpha_h)\rfloor, \qquad R_h := \lfloor A_{n,h}(\beta_h)\rfloor.$$

Then, after increasing n depending on H,

$$0 \le L_{H-1} \le R_{H-1} \le L_{H-2} \le R_{H-2} \le \cdots \le L_1 \le R_1 < K,$$

and

$$C_{H,n}(X) := \begin{cases} 0, & 2 \le X < n^{\alpha_{H-1}}, \\ \lfloor A_{n,h}(\log_n X) \rfloor, & n^{\alpha_h} \le X \le n^{\beta_h} \quad (1 \le h \le H-1), \\ R_h, & n^{\beta_h} < X < n^{\alpha_{h-1}} \quad (2 \le h \le H-1), \\ R_1, & n^{\beta_1} < X < n, \\ K, & X = n. \end{cases}$$

At each jump endpoint the value is the value on the right-hand piece; in particular the first point at which a level j is reached satisfies $C_{H,n}(b_j) \ge j$ for the generalized inverse b_j of Lemma 7.5. The endpoint algebra is the same one formalized in `Envelope.lean`*, and the reciprocal profile at the flat endpoints is mirrored by the zero-sorry lemmas in* `FlatMass.lean`*.*

Proof. By Proposition 7.3, there is a function $\xi_H(n) \to 0$ such that for every large enough n, every $1 \le h \le H-1$, and every $u \in I_h$,

$$S_K(n^u) \ge \big(1 - \xi_H(n)\big) \frac{n^u}{(h+1)u \log n}.$$

For fixed H and all sufficiently large n, the functions $A_{n,h}$ are strictly increasing on I_h, since

$$A'_{n,h}(u) = \big(1 - \xi_H(n)\big) \frac{n^u}{(h+1)u} \left(1 - \frac{1}{u \log n}\right),$$

and $u \ge \alpha_{H-1} > 1/H$.

The definitions of L_h, R_h give $L_h \le R_h$. Moreover, for $2 \le h \le H-1$,

$$\frac{A_{n,h-1}(\alpha_{h-1})}{A_{n,h}(\beta_h)} = n^{2\tau_H} \cdot \frac{(h+1)(1/h - \tau_H)}{h(1/h + \tau_H)} \to \infty,$$

so $R_h \le L_{h-1}$ for all large n. Also

$$R_1 \ll \frac{n^{1-\tau_H}}{\log n} = o(K),$$

again for large n. Thus the displayed order of endpoint levels holds.

The function $C_{H,n}$ is nondecreasing on each genuine block, constant on each gap, and nondecreasing across adjacent pieces by $R_h \le L_{h-1}$. Its endpoint convention makes it right-continuous. Domination by S_K is checked piece by piece: on a genuine block it is exactly the local-density lower bound, on a gap it follows from $C_{H,n}(X) = R_h \le S_K(n^{\beta_h}) \le S_K(X)$, and at the endpoints below $n^{\alpha_{H-1}}$ and at n it is immediate.

Finally, because $C_{H,n}$ is a finite right-continuous step function whose endpoint values are included on the right-hand pieces, its first hitting point for any level j is one of the displayed endpoints or a genuine-block point at which $\lfloor A_{n,h}(\log_n X) \rfloor \ge j$. Hence the first hitting point b_j satisfies $C_{H,n}(b_j) \ge j$. □

Lemma 7.5 (Generalized inverse comparison sequence)**.** *Let $C_{H,n}$ be the envelope of Lemma 7.4, and define*

$$b_j := \inf\{X \in [2,n] : C_{H,n}(X) \ge j\} \qquad (1 \le j \le K).$$

Then $b_1 \le \cdots \le b_K$, and Shortener's prime prefix satisfies

$$q_j \le b_j \qquad (1 \le j \le K).$$

Equivalently, the inverse is given by the endpoint-and-block formula

$$b_j = \begin{cases} n^{\alpha_{H-1}}, & 1 \le j \le L_{H-1}, \\ n^{u_{h,j}}, & L_h < j \le R_h, \ \textit{where } A_{n,h}(u_{h,j}) = j, \\ n^{\alpha_{h-1}}, & R_h < j \le L_{h-1} \quad (2 \le h \le H-1), \\ n, & R_1 < j \le K. \end{cases}$$

The abstract cutoff comparison is formalized without sorries as `cutoff_le_of_lowerProfile` *in* `Round15Bonferroni4/Inversion.lean`*.*

Proof. Monotonicity of $C_{H,n}$ gives $b_1 \le \cdots \le b_K$, and the displayed formula is just the first-hitting description of the step function. By Lemma 7.4,

$$S_K(b_j) \ge C_{H,n}(b_j) \ge j.$$

Thus at least j of the primes $q_1, \dots, q_K$ are at most b_j, so the jth one satisfies $q_j \le b_j$. □

Lemma 7.6 (Flat-block mass and repeated indices)**.** *With b_j, u_j, w_j defined by Lemma 7.5, set*

$$u_j := \log_n b_j, \qquad w_j := \frac{1}{b_j} = n^{-u_j},$$

and

$$\mu_{H,n} := \sum_{j=1}^{K} w_j \, \delta_{u_j}.$$

If $\mu_{H,n} = \nu_{H,n} + \lambda_{H,n}$ is split into genuine-block and flat filler-block contributions, then

$$|\lambda_{H,n}| = O_H(1/\log n) = o_H(1).$$

Moreover $\mu_{H,n}([0,1]) = O_H(1)$ and, for each fixed $1 \le r \le 4$,

$$0 \le \int_{\Sigma_r} d\mu_{H,n}^{\otimes r} - r!\, T_r^{(b)}(n) \le \binom{r}{2} \big(\max_j w_j\big) \mu_{H,n}([0,1])^{r-1} = o_H(1),$$

where $\Sigma_r = \{(u_1, \dots, u_r) : u_1 + \cdots + u_r \le 1\}$. The reciprocal mass estimates are the paper-level versions of `bottomFlatMass_eq` *and* `gapFlatMass_le` *in* `Round15Bonferroni4/FlatMass.lean`*.*

Proof. For the bottom flat block,

$$|\lambda_{H,n}^{\text{bot}}| = L_{H-1} n^{-\alpha_{H-1}} \le A_{n,H-1}(\alpha_{H-1}) n^{-\alpha_{H-1}} = \frac{1 - \xi_H(n)}{H\alpha_{H-1} \log n} = O_H\left(\frac{1}{\log n}\right).$$

For a gap block between I_h and I_{h-1}, where $2 \le h \le H-1$,

$$|\lambda_{H,n}^{(h)}| = (L_{h-1} - R_h) n^{-\alpha_{h-1}} \le L_{h-1} n^{-\alpha_{h-1}} \le \frac{1 - \xi_H(n)}{h\, \alpha_{h-1} \log n} = O_H\left(\frac{1}{\log n}\right).$$

For the top block,

$$|\lambda_{H,n}^{\text{top}}| = (K - R_1) n^{-1} \le \frac{K}{n} = O\left(\frac{1}{\log n}\right).$$

There are only $O_H(1)$ filler blocks, so $|\lambda_{H,n}| = O_H(1/\log n) = o_H(1)$.

The same Stieltjes calculation used below with $f \equiv 1$ gives $\nu_{H,n}([0,1]) = O_H(1)$, and the flat contribution is $o_H(1)$; hence $\mu_{H,n}([0,1]) = O_H(1)$. The diagonal or repeated-index part of the ordered r-fold product is bounded by

$$\binom{r}{2} \big(\max_j w_j\big) \mu_{H,n}([0,1])^{r-1}.$$

Since $\max_j w_j \le n^{-\alpha_{H-1}} \to 0$ for fixed H, this is $o_H(1)$. □

Lemma 7.7 (Log-scale measure convergence). *Let $G_H = \bigcup_{h=1}^{H-1} I_h$. For every $f \in C^1([0,1])$,*

$$\int f \, d\mu_{H,n} = \int_{G_H} f(u)\rho(u)\,du + o_H(1).$$

Consequently, for every fixed $1 \le r \le 4$,

$$\int_{\Sigma_r} d\mu_{H,n}^{\otimes r} \to \int_{\substack{u_1+\cdots+u_r\le 1\\ u_i\in G_H}} \prod_{i=1}^r \rho(u_i)\,du_1\cdots du_r.$$

The atoms from flat blocks have vanishing total mass by Lemma 7.6; the remaining discontinuities are only the finitely many coordinate breakpoint hyperplanes and the hyperplane $u_1+\cdots+u_r = 1$, which have zero limiting product measure.

Proof. Fix $1 \le h \le H-1$. On the block $L_h < j \le R_h$, the value $u_j =: u_{h,j} \in I_h$ is defined by $A_{n,h}(u_{h,j}) = j$. Therefore, for every $f \in C^1([0,1])$,

$$\int f\,d\nu_{H,n}^{(h)} = \sum_{j=L_h+1}^{R_h} f(u_{h,j})n^{-u_{h,j}} = \int_{\alpha_h}^{\beta_h} f(u)n^{-u}\,d\lfloor A_{n,h}(u)\rfloor.$$

Let $R_{n,h}(u) := \lfloor A_{n,h}(u)\rfloor - A_{n,h}(u)$, so $|R_{n,h}(u)| \le 1$. Stieltjes integration by parts with $g(u) := f(u)n^{-u}$ gives

$$\int_{\alpha_h}^{\beta_h} g(u)\,d\lfloor A_{n,h}(u)\rfloor = \int_{\alpha_h}^{\beta_h} g(u)\,dA_{n,h}(u) + E_{n,h}(f),$$

where

$$|E_{n,h}(f)| \le 2\|g\|_\infty + \int_{\alpha_h}^{\beta_h} |g'(u)|\,du = o_H(1),$$

because $g'(u) = \big(f'(u) - (\log n)f(u)\big)n^{-u}$ and $\alpha_h \ge \alpha_{H-1} > 1/H$. On the other hand,

$$n^{-u}A'_{n,h}(u) = \big(1-\xi_H(n)\big)\rho(u)\left(1 - \frac{1}{u\log n}\right) \qquad (u \in I_h),$$

because $\rho(u) = 1/((h+1)u)$ on I_h. Hence

$$\int_{\alpha_h}^{\beta_h} f(u)n^{-u}\,dA_{n,h}(u) = \int_{I_h} f(u)\rho(u)\,du + o_H(1),$$

and therefore

$$\int f\,d\nu_{H,n}^{(h)} = \int_{I_h} f(u)\rho(u)\,du + o_H(1).$$

Summing over h and using Lemma 7.6 gives the displayed convergence for C^1 test functions, and hence for continuous test functions by uniform approximation.

The tensor-product convergence follows for continuous functions on $[0,1]^r$. The indicator of Σ_r restricted to G_H^r is Riemann integrable: the only discontinuities are the finitely many coordinate breakpoints and $u_1+\cdots+u_r = 1$, each of zero Lebesgue measure against $\prod_i \rho(u_i)\,du_i$. Approximating this indicator from above and below by continuous functions gives the stated limit. ☐

Proposition 7.8 (Envelope and inversion). *Fix $H \ge 2$, and define $\tau_H, \alpha_h, \beta_h, I_h$ as in Proposition 7.3. Put*

$$G_H := \bigcup_{h=1}^{H-1} I_h,$$

and for $1 \le r \le 4$ define

$$J_r^{(H)} := \frac{1}{r!}\int_{\substack{u_1+\cdots+u_r\le 1\\ u_i\in G_H}} \prod_{i=1}^r \rho(u_i)\,du_1\cdots du_r.$$

Then for every sufficiently large n there exists a monotone function $C_{H,n} : [2,n] \to \{0,1,\dots,K\}$ such that:

(1) $C_{H,n}(X) \le S_K(X)$ *for all* $2 \le X \le n$*;*

(2) *if*

$$b_j := \inf\{X \in [2,n] : C_{H,n}(X) \ge j\},$$

then $q_j \le b_j$ *for every* $1 \le j \le K$*;*

(3) *for every fixed* $r \in \{1,2,3,4\}$,

$$T_r^{(b)}(n) = J_r^{(H)} + o_H(1) \qquad (n \to \infty);$$

(4) *for every fixed* $r \in \{1,2,3,4\}$,

$$0 \le J_r - J_r^{(H)} \le \frac{5r}{4H\, r!}.$$

In particular, there exists an integer-valued function $H_*(n) \to \infty$ and a corresponding choice of comparison sequences

$$b_1^{(n)} \le \dots \le b_K^{(n)}$$

obtained from the above construction with $H = H_*(n)$ such that

$$T_r^{(b^{(n)})}(n) = J_r + o(1) \qquad (1 \le r \le 4).$$

Proof. Lemma 7.4 gives the dominated envelope, and Lemma 7.5 gives $q_j \le b_j$. By Lemma 7.6, repeated ordered tuples contribute only $o_H(1)$ to the r-fold product integral. Combining this with Lemma 7.7 gives

$$T_r^{(b)}(n) = J_r^{(H)} + o_H(1).$$

It remains to control the excised breakpoint region. Since $0 < \rho(u) < 1$ on $(0,1]$, we have

$$0 \le J_r - J_r^{(H)} \le \frac{1}{r!}\, \lambda_r\big(\{(u_1,\dots,u_r) \in [0,1]^r : \exists i,\ u_i \notin G_H\}\big),$$

where λ_r denotes r-dimensional Lebesgue measure. Moreover,

$$|(0,1] \setminus G_H| = \frac{1}{H} + 2(H-1)\tau_H \le \frac{1}{H} + \frac{1}{4H} = \frac{5}{4H}.$$

By the union bound in $[0,1]^r$,

$$\lambda_r\big(\{(u_1,\dots,u_r) \in [0,1]^r : \exists i,\ u_i \notin G_H\}\big) \le r\,|(0,1] \setminus G_H| \le \frac{5r}{4H}.$$

This proves the fixed-H assertions.

To obtain a single comparison sequence for each n, choose N_H so large that for all $n \ge N_H$ the above construction is valid and, for every $1 \le r \le 4$,

$$\big|T_r^{(b)}(n) - J_r^{(H)}\big| \le \frac{1}{H}.$$

After enlarging if necessary, assume N_H is increasing in H. Define

$$H_*(n) := \max\{H \ge 2 : N_H \le n\},$$

and let $b_1^{(n)} \le \dots \le b_K^{(n)}$ be the comparison sequence produced with parameter $H = H_*(n)$. Then $H_*(n) \to \infty$, and for every $1 \le r \le 4$,

$$\big|T_r^{(b^{(n)})}(n) - J_r\big| \le \big|T_r^{(b^{(n)})}(n) - J_r^{(H_*(n))}\big| + \big|J_r^{(H_*(n))} - J_r\big| \le \frac{1}{H_*(n)} + \frac{5r}{4H_*(n)r!},$$

which tends to 0. □

7.4. **The prime-sequence Bonferroni comparison theorem.** The comparison theorem must be stated for *prime sequences*, not for arbitrary reals; the correct statement is the following.

Lemma 7.9 (Odd-sieve monotonicity under prime upgrades). *Let $q_1 < \cdots < q_K$ and $p_1 < \cdots < p_K$ be odd primes with $q_j \le p_j$ for every j. If $N(r_1, \ldots, r_K)$ denotes the number of odd integers $m \le n$ not divisible by any of the primes $r_1, \ldots, r_K$, then*

$$N(q_1, \ldots, q_K) \le N(p_1, \ldots, p_K).$$

Proof. Replace the primes one at a time, in descending order of the index. Suppose the primes other than the kth one form a set R, and let $q \le p$ be the two possible kth primes, neither lying in R. Let $A_R(t)$ be the number of odd integers at most t which are divisible by no prime in R. Since q and p are distinct from every prime in R,

$$N(R \cup \{q\}) = A_R(n) - A_R(n/q), \qquad N(R \cup \{p\}) = A_R(n) - A_R(n/p).$$

The function $A_R(t)$ is nondecreasing, and $q \le p$, so $A_R(n/q) \ge A_R(n/p)$. Hence replacing q by p can only increase the survivor count. Descending replacement preserves distinctness and the sorted structure, because after the indices $> k$ have been upgraded we still have $q_{k-1} < q_k \le p_k < p_{k+1}$, with the obvious endpoint omissions. □

Theorem 7.10 (Prime-sequence Bonferroni–4 comparison). *Let $q_1 < \cdots < q_K$ be the actual odd primes played by Shortener during the prime prefix, and let $p_1 < \cdots < p_K$ be any increasing sequence of odd primes with*

$$q_j \le p_j \qquad (1 \le j \le K).$$

Assume that for $1 \le r \le 4$,

$$T_r^{(p)}(n) = \Lambda_r + o(1).$$

Then

$$L(n) \le \frac{n}{2} \left(1 - \Lambda_1 + \Lambda_2 - \Lambda_3 + \Lambda_4 + o(1)\right).$$

Proof. Let A be the final played set. If the game ends before Shortener completes the prefix of length K, then $L(n) = o(n)$ and there is nothing to prove, so assume $q_1, \ldots, q_K$ were all played.

Define the odd-part map

$$\varphi(x) := \frac{x}{2^{v_2(x)}}.$$

If $\varphi(x) = \varphi(y) = m$, then $x = 2^a m$ and $y = 2^b m$. If $a \ne b$, one of x and y divides the other, contradicting that A is an antichain. So φ is injective on A.

Since each $q_j \in A$, no other element of A can be divisible by any q_j. Because q_j is odd, divisibility by q_j is preserved under the odd-part map. Hence

$$\varphi(A) \subseteq \{m \le n : m \text{ odd and } q_j \nmid m \text{ for all } j\} \cup \{q_1, \ldots, q_K\}.$$

Therefore

$$L(n) \le N(q_1, \ldots, q_K) + K = N(q_1, \ldots, q_K) + o(n),$$

where $N(q_1, \ldots, q_K)$ counts odd integers up to n avoiding all q_j.

Now replace the actual sequence q_j by the larger prime sequence p_j. Lemma 7.9 gives

$$N(q_1, \ldots, q_K) \le N(p_1, \ldots, p_K).$$

Applying fourth-order Bonferroni to the divisibility events $p_j \mid m$ over odd integers $m \le n$ gives

$$\begin{aligned} N(p_1, \ldots, p_K) \le \frac{n}{2} &\left(1 - T_1^{(p)}(n) + T_2^{(p)}(n) - T_3^{(p)}(n) + T_4^{(p)}(n)\right) \\ &+ O(1 + D_1 + D_2 + D_3 + D_4), \end{aligned}$$

where

$$D_r := \#\{j_1 < \cdots < j_r : p_{j_1} \cdots p_{j_r} \le n\}.$$

Indeed, for a fixed r-tuple $J = \{j_1 < \cdots < j_r\}$ with $P_J := \prod_{j \in J} p_j \le n$, the number of odd $m \le n$ divisible by P_J is

$$\frac{n}{2P_J} + O(1),$$

because P_J is odd. If $P_J > n$ the intersection is empty, so the total floor error in the first four Bonferroni sums is $O(D_1 + D_2 + D_3 + D_4)$. The additional $O(1)$ accounts for replacing the exact number of odd integers at most n by $n/2$.

By Landau's theorem on the count of r-almost-primes [HW08, Theorem 437],

$$\#\{m \le n : m \text{ squarefree and } \omega(m) = r\} = O_r\left(\frac{n(\log\log n)^{r-1}}{\log n}\right).$$

Every product counted by D_r is such a squarefree integer, and therefore

$$D_r = O_r\left(\frac{n(\log\log n)^{r-1}}{\log n}\right) = o(n) \qquad (1 \le r \le 4),$$

so the error is $o(n)$. Substituting $T_r^{(p)}(n) = \Lambda_r + o(1)$ yields the stated inequality. □

7.5. **The prime-rounding bridge.**

Lemma 7.11 (Fixed-ratio cell reciprocal mass)**.** *Fix $H \ge 2$, and let $b_1, \dots, b_K$ be the comparison sequence obtained from Proposition 7.8 with this fixed H. For every fixed $c > 1$, uniformly in $Y > 0$,*

$$\sum_{\substack{1 \le j \le K \\ Y \le b_j < cY}} \frac{1}{b_j} = O_{H,c}\left(\frac{1}{\log n}\right).$$

Proof. If the cell $[Y, cY)$ lies inside a genuine block $[n^{\alpha_h}, n^{\beta_h}]$, then $B_{H,n}(X) := \#\{j : b_j \le X\}$ agrees with $\lfloor A_{n,h}(\log_n X) \rfloor$ on the block. As in the proof of Proposition 7.8,

$$\frac{d}{dX} A_{n,h}(\log_n X) = \frac{1 - \xi_H(n)}{(h+1)\log X}\left(1 - \frac{1}{\log X}\right) \ll_H \frac{1}{\log n},$$

because every genuine block has $X \ge n^{\alpha_{H-1}} > n^{1/H}$. Hence

$$\#\{j : Y \le b_j < cY,\ j \text{ genuine}\} \ll_{H,c} \frac{Y}{\log n} + 1.$$

Dividing by Y gives a reciprocal contribution $O_{H,c}(1/\log n)$; the endpoint $+1$ is harmless since $Y \ge n^{1/H}$ in a genuine block.

A fixed-ratio cell can meet at most one breakpoint for all large n, because successive breakpoints in the envelope are separated by a factor which is a positive power of n. If the cell is adjacent to a breakpoint, split its genuine part into the at most two adjacent genuine pieces and apply the preceding estimate to each. The only additional contribution is the flat atom at the breakpoint. At $A_h := n^{\alpha_h}$ the flat multiplicity is at most L_h, so

$$\frac{L_h}{A_h} \le \frac{A_{n,h}(\alpha_h)}{n^{\alpha_h}} = O_H\left(\frac{1}{\log n}\right).$$

The top flat block at $b_j = n$ contributes at most

$$\frac{K}{n} = O\left(\frac{1}{\log n}\right).$$

Cells lying wholly in gaps have no genuine contribution and only the same possible flat-atom contribution. These cases cover all Y. □

Lemma 7.12 (Product-strip reciprocal-mass estimate). *Fix $H \geq 2$, fixed real numbers $C > 1$ and $c > 1$, and an integer $r \geq 1$. Let $b_1, \dots, b_K$ be the comparison sequence obtained from Proposition 7.8 with this fixed H. Then, for all sufficiently large n,*

$$\sum_{\substack{j_1 < \cdots < j_r \\ n/C < b_{j_1} \cdots b_{j_r} \leq Cn}} \frac{1}{b_{j_1} \cdots b_{j_r}} = O_{H,C,c,r}\left(\frac{1}{\log n}\right).$$

Proof. Partition $(1, n]$ into logarithmic cells

$$\mathcal{C}_m := [c^m, c^{m+1}) \qquad (0 \leq m \leq \lceil \log_c n \rceil).$$

There are $O_c(\log n)$ possible cells in each coordinate. In logarithmic coordinates, the condition

$$n/C < x_1 \cdots x_r \leq Cn$$

is a strip of width $O_{C,c,r}(1)$ cell-widths around the hyperplane $y_1 + \cdots + y_r = \log n$. Equivalently, if $x_i \in \mathcal{C}_{m_i}$ and the product cell meets this strip, then

$$\log_c n - O_{C,c,r}(1) \leq m_1 + \cdots + m_r \leq \log_c n + O_{C,c,r}(1).$$

Thus, after any $r - 1$ cell indices have been chosen, the last cell index is restricted to $O_{C,c,r}(1)$ possibilities. The number of relevant cell patterns is therefore $O_{C,c,r}((\log n)^{r-1})$.

For each cell pattern, the contribution of distinct tuples lying in those cells is bounded by the product of the corresponding cell reciprocal masses. By Lemma 7.11, each such cell has reciprocal mass $O_{H,c}(1/\log n)$, including cells containing flat atoms or the top atom $b_j = n$. Hence every relevant r-cell pattern contributes $O_{H,c,r}((\log n)^{-r})$. Multiplying by the number of relevant patterns gives

$$O_{H,C,c,r}((\log n)^{r-1})\, O_{H,c,r}((\log n)^{-r}) = O_{H,C,c,r}\left(\frac{1}{\log n}\right),$$

as claimed. □

It remains to convert the real comparison sequence b_j from Proposition 7.8 into an actual increasing odd-prime sequence. The naive claim

$$p_j/b_j = 1 + O(1/\log b_j)$$

uniformly in j is *not* justified by the prime number theorem alone and must not be used. The obstruction is concrete: the top flat block contributes up to $K \leq n/(2\log n)$ indices with $b_j = n$, yet every window $[n, (1 + o(1))n]$ contains only $o(n/\log n)$ odd primes, so no uniform $p_j/b_j = 1 + o(1)$ can supply the top block. The correct bridge is the half-density construction of Theorem 7.20, which assigns the top flat block separately to $[n, 2n]$—whose $(1 + o(1))n/\log n$ odd primes exceed the demand K by a factor of two—and uses queued coarse-bin matching on the genuine blocks, where the envelope bound $\frac{1}{(h+1)\log X} \leq \frac{1}{2\log X}$ on the local b-density supplies slack against the odd-prime density $1/\log X$.

Lemma 7.13 (Bin demand estimate). *Fix $H \geq 2$ and a genuine block $[A_h, B_h]$, where $A_h = n^{\alpha_h}$ and $B_h = n^{\beta_h}$. For a fixed $a > 1$, partition the block into bins*

$$I_k = [A_h a^k, A_h a^{k+1})$$

with the last bin truncated at B_h, and let

$$d_k := \#\{j : b_j \in I_k,\ j \text{ lies in the genuine } h\text{-block}\}.$$

Then, for every fixed $c > 1$, uniformly for $X \geq n^{1/H}$ with $[X, cX] \subseteq [A_h, B_h]$,

$$\#\{j : X \leq b_j < cX,\ j \text{ genuine in block } h\} \leq (1 + o_{H,c}(1))\frac{(c-1)X}{(h+1)\log X} + O(1). \tag{7.2}$$

If f_h denotes the multiplicity of the flat atom at the left endpoint A_h, then

$$f_h \le (1 + o_H(1)) \frac{A_h}{(h+1)\log A_h}. \tag{7.3}$$

Proof. On a genuine block, the inverse sequence has $B_{H,n}(Y) = \lfloor A_{n,h}(\log_n Y)\rfloor$, and

$$\frac{d}{dY} A_{n,h}(\log_n Y) = \frac{1 - \xi_H(n)}{(h+1)\log Y}\left(1 - \frac{1}{\log Y}\right).$$

Integrating over $[X, cX]$ gives (7.2); the $O(1)$ endpoint error is negligible at these scales. The flat atom satisfies $f_h \le L_h \le A_{n,h}(\alpha_h)$, which gives (7.3). ☐

Lemma 7.14 (Bin prime supply estimate)**.** *With $A = A_h$ fixed as above, let*

$$P_k := \#\{p \text{ odd prime} : Aa^k \le p < Aa^{k+1}\}.$$

Then the Prime Number Theorem in fixed multiplicative intervals [HW08, Theorem 6] *gives, uniformly in k in this fixed block,*

$$P_k = (1 + o_{H,a}(1)) \frac{(a-1)Aa^k}{\log(Aa^k)}. \tag{7.4}$$

Proof. For fixed H and a, all bins in a genuine block have endpoints tending to infinity and fixed ratio a. Applying the fixed-ratio PNT to $[Aa^k, Aa^{k+1})$ gives the displayed estimate. The possible prime 2 is irrelevant because these intervals are eventually above 2. ☐

Lemma 7.15 (Interval-order Hall suffix inequality)**.** *Fix $a > 1$. Choose $\eta > 0$ with $2a(1-\eta)/(1+\eta) > 1$, and then choose $s = s(a)$ so that*

$$(1-\eta)(a^{s+1} - 1) > (1+\eta)\frac{a^s}{2}. \tag{7.5}$$

For all sufficiently large n, the flat atom and the first s genuine bins in each genuine block satisfy the Hall suffix inequalities

$$\sum_{\ell=t}^{s} P_\ell \ge \sum_{k=t-1}^{s-1} d_k \qquad (1 \le t \le s),$$

and

$$\sum_{\ell=0}^{s} P_\ell \ge f_h + \sum_{k=0}^{s-1} d_k.$$

Proof. Throughout this proof $A = A_h = n^{\alpha_h} \to \infty$ while $s = s(a)$ is fixed, so

$$\log(Aa^j) = \log A + j \log a = (1 + o_{H,a}(1)) \log A$$

uniformly for $0 \le j \le s+1$. The resulting $o(1)$ multiplicative errors in the denominators below are absorbed into the $(1 \pm \eta)$ margins, all valid for all large n.

For $1 \le t \le s$, Lemma 7.14 gives

$$\sum_{\ell=t}^{s} P_\ell \ge (1-\eta)\frac{A(a^{s+1} - a^t)}{\log A} = (1-\eta)\frac{aA(a^s - a^{t-1})}{\log A},$$

for all large n, while Lemma 7.13 gives

$$\sum_{k=t-1}^{s-1} d_k \le (1+\eta)\frac{A(a^s - a^{t-1})}{2\log A}.$$

The suffix supply-to-demand ratio is therefore at least $2a(1-\eta)/(1+\eta) > 1$.

For $t = 0$, the demand and flat-atom estimates give

$$f_h + \sum_{k=0}^{s-1} d_k \le (1+\eta)\frac{Aa^s}{2\log A},$$

and the prime-bin estimate gives

$$\sum_{\ell=0}^{s} P_\ell \ge (1-\eta)\frac{A(a^{s+1}-1)}{\log A}.$$

The choice (7.5) proves the $t = 0$ inequality. □

Lemma 7.16 (Monotone assignment)**.** *Under the hypotheses of Lemma 7.15, the flat atom and all genuine-bin demands in a fixed genuine block can be assigned to odd prime supply bins so that the assigned primes are strictly increasing and at least as large as their corresponding b_j. The endpoint flat atom of the block is queued with release index 0; a genuine demand from source bin I_k is released only at prime bin $k+1$. The top flat block $b_j = n$ is assigned after all genuine blocks, in order, to odd primes in $[n, 2n]$.*

Proof. The suffix inequalities of Lemma 7.15 are exactly the Hall inequalities for this interval-order matching. Equivalently, process prime bins in increasing order and assign each prime to the oldest unassigned released demand. If this greedy algorithm first failed before or at bin s, the unassigned demands with release index at least the failure bin would violate the corresponding suffix inequality.

For each later full source bin I_k, $k \ge s$, assign its d_k elements to primes in the next prime bin $[Aa^{k+1}, Aa^{k+2})$. By Lemmas 7.13 and 7.14,

$$d_k \le (1+o_{H,a}(1))\frac{(a-1)Aa^k}{(h+1)\log(Aa^k)} < (1-o_{H,a}(1))\frac{(a-1)Aa^{k+1}}{\log(Aa^{k+1})} \le P_{k+1},$$

using $1/(h+1) \le 1/2 < a$. For the final truncated source bin $[X, B_h] \subseteq [Aa^{k^\star}, Aa^{k^\star+1})$ (with $X = Aa^{k^\star}$ and $B_h < Aa^{k^\star+1}$), the same derivative bound used in Lemma 7.13 over $[X, B_h]$ gives

$$d_{k^\star} \le (1+o_{H,a}(1))\frac{B_h - X}{(h+1)\log X} + O(1) \le (1+o_{H,a}(1))\frac{(a-1)X}{(h+1)\log X} + O(1) < P_{k^\star+1},$$

since $B_h/X \le a$ and the next prime bin above B_h is contained in $[Aa^{k^\star+1}, Aa^{k^\star+2})$. Thus the truncated-bin assignment proceeds by the same greedy rule.

The greedy rule, together with processing bins in increasing order and taking the smallest available prime in each bin, makes the assigned sequence strictly increasing within each block. Assignments from different genuine blocks do not overlap, since the largest prime used for block h is at most $a^2 B_h$, whereas the next larger block begins at

$$A_{h-1} = n^{1/h+\tau_H} = n^{2\tau_H} B_h,$$

so $a^2 B_h < A_{h-1}$ for large n; also $a^2 B_1 = o(n)$. Finally $[n, 2n]$ contains $(1+o(1))n/\log n$ odd primes, more than enough for the top flat block of size at most $K \le n/(2\log n)$. □

Lemma 7.17 (Exceptional mass and cutoff discrepancy)**.** *Let E be the union of the flat atoms, the first $s(a)$ genuine bins of every genuine block, and the top flat block $b_j = n$. Then*

$$\sum_{j\in E} \frac{1}{b_j} = O_{H,a}\left(\frac{1}{\log n}\right), \tag{7.6}$$

and the same bound holds with p_j in place of b_j. For every nonexceptional index,

$$b_j \le p_j \le a^2 b_j \le \lambda b_j. \tag{7.7}$$

Consequently, for every fixed $1 \le r \le 4$,

$$T_r^{(p)}(n) = T_r^{(b)}(n) + O_{H,r}(\lambda - 1) + o_{H,a,s,\lambda}(1),$$

where the $o(1)$ error depends on the block count H, the fixed bin ratio a, the coarse-bin prefix count $s(a)$ of Lemma 7.17 above, and the queue parameter λ (all held fixed before the diagonalization of Proposition 7.19 lets $\lambda \downarrow 1$ and $H \to \infty$).

Proof. By Lemma 7.11, each fixed multiplicative bin has b-reciprocal mass $O_{H,a}(1/\log n)$, and there are only $O_{H,a}(1)$ exceptional bins. The top flat block contributes $K/n = O(1/\log n)$, proving (7.6). Since every assigned prime is at least the corresponding b_j, the same bound holds with p_j in place of b_j.

For a nonexceptional index, the construction of Lemma 7.16 puts p_j in the next prime bin above the bin containing b_j, so $b_j \le p_j \le a^2 b_j \le \lambda b_j$. All r-tuples containing at least one exceptional index contribute $O_{H,r,a}(1/\log n) = o_{H,\lambda}(1)$ to either moment sum. For a nonexceptional tuple I, write $b_I = \prod_{j\in I} b_j$ and $p_I = \prod_{j\in I} p_j$. Then

$$b_I \le p_I \le \lambda^r b_I, \qquad \lambda^{-r}\frac{1}{b_I} \le \frac{1}{p_I} \le \frac{1}{b_I}.$$

On tuples counted by both cutoffs, this changes the total reciprocal weight by $O_{H,r}(\lambda-1)$, since the r-fold reciprocal mass is $O_H(1)$. The only cutoff discrepancy comes from tuples with $n/\lambda^r < b_I \le n$, whose total b-reciprocal contribution is $O_{H,\lambda,a,r}(1/\log n)$ by Lemma 7.12 with $C = \lambda^r$ and $c = a$. ☐

Proposition 7.18 (Fixed-parameter queued prime rounding). *Fix $H \ge 2$ and $\lambda > 1$, and choose $a > 1$ with $a^2 \le \lambda$. Let $b_1, \dots, b_K$ be the comparison sequence obtained from Proposition 7.8 with this fixed H. Then there is an increasing sequence of odd primes $p_1 < \cdots < p_K$ such that*

$$q_j \le b_j \le p_j \qquad (1 \le j \le K),$$

and, for every fixed $1 \le r \le 4$,

$$T_r^{(p)}(n) = T_r^{(b)}(n) + O_{H,r}(\lambda - 1) + o_{H,\lambda}(1). \tag{7.1}$$

Proof. Write $A_h = n^{\alpha_h}$ and $B_h = n^{\beta_h}$. In each genuine block, Lemmas 7.13 and 7.14 provide the demand and supply estimates for the fixed multiplicative bins. Choose η and $s = s(a)$ as in Lemma 7.15. The suffix inequalities there give the finite interval-order Hall matching for the endpoint flat atom and the first s genuine bins, while Lemma 7.16 assigns all later genuine bins, keeps the assignments from different blocks ordered, and sends the top flat block to primes in $[n, 2n]$. Thus we obtain a strictly increasing odd-prime sequence with $b_j \le p_j$ for every j. Together with Proposition 7.8, this gives $q_j \le b_j \le p_j$.

The moment comparison is exactly Lemma 7.17, which proves (7.1) for each fixed $1 \le r \le 4$. ☐

Proposition 7.19 (Diagonal parameter selection for the prime bridge). *Fix $0 < \varepsilon < 1$. There is an integer-valued function $H(n) \to \infty$ and, for each n, comparison sequences*

$$b_1^{(n)} \le \cdots \le b_K^{(n)}, \qquad p_1^{(n)} < \cdots < p_K^{(n)}$$

with the $p_j^{(n)}$ odd primes, such that

$$q_j \le b_j^{(n)} \le p_j^{(n)} \qquad (1 \le j \le K),$$

and, for every $1 \le r \le 4$,

$$T_r^{(b^{(n)})}(n) = J_r + o(1), \qquad T_r^{(p^{(n)})}(n) = J_r + o(1).$$

Proof. Choose $H_m \to \infty$ increasing so that the excision estimate in Proposition 7.8 gives

$$|J_r^{(H_m)} - J_r| \le \frac{1}{4m} \qquad (1 \le r \le 4). \tag{7.8}$$

For each fixed H_m and r, choose a constant $C_{H_m,r}$ for the $O_{H_m,r}(\lambda-1)$ term in (7.1). Choose $\lambda_m > 1$ close enough to 1 that

$$C_{H_m,r}(\lambda_m - 1) \le \frac{1}{4m} \qquad (1 \le r \le 4), \tag{7.9}$$

and then choose $a_m > 1$ with $a_m^2 \le \lambda_m$. In the queue construction with this a_m, fix $\eta_m > 0$ small enough that $2a_m(1-\eta_m)/(1+\eta_m) > 1$, and let s_m be a corresponding queue length for (7.5). Let $b^{(H_m)}$ be the comparison sequence from Proposition 7.8 with parameter H_m, and let $p^{(H_m)}$ be the prime sequence produced by Proposition 7.18 with parameters H_m, λ_m, a_m.

Choose N_m recursively increasing so that for every $n \ge N_m$, the twelve diagonal selection conditions listed in Appendix A.2 hold for the fixed parameters $H_m, \lambda_m, a_m, \eta_m, s_m$. Each condition is eventual because these parameters and the finite list of ratios are fixed while $n \to \infty$.

Define

$$H(n) := H_m \qquad \text{for } N_m \le n < N_{m+1},$$

and use the corresponding sequences $b^{(H_m)}$ and $p^{(H_m)}$. Then $H(n) \to \infty$. For $N_m \le n < N_{m+1}$, the telescope bounds are

$$\begin{aligned} |T_r^{(b^{(H(n))})}(n) - J_r| &\le |T_r^{(b^{(H_m)})}(n) - J_r^{(H_m)}| + |J_r^{(H_m)} - J_r| \\ &\le \frac{1}{4m} + \frac{1}{4m} = \frac{1}{2m}, \end{aligned}$$

and

$$\begin{aligned} |T_r^{(p^{(H(n))})}(n) - J_r| &\le |T_r^{(p^{(H_m)})}(n) - T_r^{(b^{(H_m)})}(n)| \\ &\quad + |T_r^{(b^{(H_m)})}(n) - J_r^{(H_m)}| + |J_r^{(H_m)} - J_r| \\ &\le \frac{1}{2m} + \frac{1}{4m} + \frac{1}{4m} = \frac{1}{m}. \end{aligned}$$

Thus $T_r^{(b^{(n)})}(n) = J_r + o(1)$ and $T_r^{(p^{(n)})}(n) = J_r + o(1)$ for $1 \le r \le 4$. The inequalities $q_j \le b_j^{(n)} \le p_j^{(n)}$ come from Proposition 7.18. □

Theorem 7.20 (Prime-rounding bridge). *Fix $0 < \varepsilon < 1$, and let*

$$K = \left\lfloor \frac{(1-\varepsilon)n}{2\log n} \right\rfloor.$$

There exists a choice of comparison sequences

$$b_1^{(n)} \le \cdots \le b_K^{(n)}$$

from the construction of Proposition 7.8, and there exists an increasing sequence of odd primes

$$p_1 < \cdots < p_K$$

such that

$$q_j \le p_j \qquad (1 \le j \le K)$$

and

$$T_r^{(b^{(n)})}(n) = J_r + o(1) \qquad (1 \le r \le 4).$$

Moreover the primes may be chosen so that, for each fixed $1 \le r \le 4$,

$$T_r^{(p)}(n) = T_r^{(b^{(n)})}(n) + o(1) = J_r + o(1).$$

Proof. The envelope side is Lemmas 7.4 to 7.7, assembled in Proposition 7.8. The rounding side is Lemmas 7.13 to 7.17, assembled in Proposition 7.18. Selecting the diagonal parameters by Proposition 7.19 gives the sequences and moment limits; since $q_j \le b_j^{(n)} \le p_j^{(n)}$, in particular $q_j \le p_j^{(n)}$. □

7.6. **The finite Bonferroni–4 theorem.**

Theorem 7.21. *Under the piecewise-density prime-prefix strategy,*

$$L(n) \le \left(\frac{\mathcal{W}_4}{2} + o(1)\right) n \le (0.1897123371 + o(1))n < 0.19n.$$

Proof. By Theorem 7.20, there are a comparison sequence $b_1 \le \cdots \le b_K$ and an increasing odd-prime sequence p_j with $q_j \le p_j$ such that

$$T_r^{(b)}(n) = J_r + o(1) \qquad (1 \le r \le 4).$$

The same theorem gives

$$T_r^{(p)}(n) = J_r + o(1) \qquad (1 \le r \le 4).$$

Applying Theorem 7.10 with $\Lambda_r = J_r$ gives

$$L(n) \le \frac{n}{2}(1 - J_1 + J_2 - J_3 + J_4 + o(1)) = \left(\frac{\mathcal{W}_4}{2} + o(1)\right) n.$$

By Proposition B.1,

$$\frac{\mathcal{W}_4}{2} \le 0.1897123371 < 0.19,$$

so the theorem follows. □

Remark 7.22. The corresponding formal artifact verifies the endgame reduction: once the first four moment terms are sufficiently close to the certified constants $J_1, \ldots, J_4$, the strict inequality $L(n) < 0.19n$ follows eventually. Theorem 7.20 supplies the moment convergence from the envelope sequence to actual primes, and Theorem 7.10 applies the fourth-order Bonferroni comparison to the resulting prime sequence.

8. Obstructions to Further Improvement

This section records three proof-class obstructions encountered by natural attempts to improve the upper bound. They are not used in the proof of Theorem 7.21; their purpose is to delineate which proof classes cannot by themselves resolve the asymptotic question.

8.1. **A local-relaxation obstruction.** The first obstruction says that bounded-round local relaxations can have an integrality gap at precisely the scale where one would hope for a short transversal argument. We record it for sparse ℓ-uniform families rather than for the complete hypergraph.

For an ℓ-uniform family $\mathcal{H} \subseteq \binom{P}{\ell}$, let $\tau_{\mathbb{Z}}(\mathcal{H})$ denote its integral transversal number and let $\tau_f(\mathcal{H})$ denote the value of the fractional set-cover relaxation

$$\min \sum_{u \in P} x_u \quad \text{subject to} \quad \sum_{u \in C} x_u \ge 1 \ (C \in \mathcal{H}), \qquad 0 \le x_u \le 1.$$

Proposition 8.1 (Sherali–Adams transversal barrier)**.** *Fix $0 < \alpha < 1$. For all sufficiently large N, put*

$$\ell = \lfloor N/\log N \rfloor, \qquad q = \lfloor \alpha N \rfloor, \qquad \delta = \frac{\binom{N-q}{\ell}}{\binom{N}{\ell}}.$$

For every N-point set P there is a family $\mathcal{C} \subseteq \binom{P}{\ell}$ such that every $Q \subseteq P$ with $|Q| \le q$ is disjoint from some $C \in \mathcal{C}$ and

$$|\mathcal{C}| \le \left\lceil \frac{\log \binom{N}{q} + 1}{\delta} \right\rceil.$$

Consequently

$$\tau_{\mathbb{Z}}(\mathcal{C}) \ge q + 1 > \alpha N, \qquad \tau_f(\mathcal{C}) \le \frac{N}{\ell} = (1 + o(1)) \log N.$$

Moreover, for every ℓ-uniform family $\mathcal{H} \subseteq \binom{P}{\ell}$ and every integer $0 \le r < \ell$, the level-r Sherali–Adams [SA90] *relaxation of the above transversal LP has value at most*

$$\frac{N}{\ell - r}.$$

Proof. Choose

$$s = \left\lceil \frac{\log \binom{N}{q} + 1}{\delta} \right\rceil$$

independent uniformly random ℓ-subsets of P. For a fixed q-set $Q \subseteq P$, a random ℓ-set avoids Q with probability δ. Thus the probability that no chosen set avoids this Q is at most $\exp(-s\delta) \le e^{-1}/\binom{N}{q}$. A union bound over all q-sets gives a positive probability that every q-set is avoided by at least one chosen ℓ-set. Removing duplicate choices gives the desired family $\mathcal{C}$. If $Q' \subseteq P$ has $|Q'| \le q$, extend it to a q-set Q; any $C \in \mathcal{C}$ disjoint from Q is also disjoint from Q'.

The avoidance property forces every transversal to have size at least $q+1$, proving $\tau_{\mathbb{Z}}(\mathcal{C}) > \alpha N$. The uniform fractional assignment $x_u = 1/\ell$ satisfies every covering constraint, so $\tau_f(\mathcal{C}) \le N/\ell$.

It remains to prove the Sherali–Adams bound. Let $\mathcal{H} \subseteq \binom{P}{\ell}$ be arbitrary and fix $0 \le r < \ell$. Set

$$p = \frac{1}{\ell - r}.$$

For a covering row $C \in \mathcal{H}$ and disjoint sets $I, J \subseteq P$ with $|I \cup J| \le r$, write

$$L_{I,J}(x) = \prod_{i \in I} x_i \prod_{j \in J} (1 - x_j).$$

In the row-wise lifted formulation, introduce the literal-moment value

$$y_{C,I,J} := \mathbb{P}(I \subseteq Y,\ J \cap Y = \emptyset) = p^{|I|}(1-p)^{|J|},$$

where Y is a product Bernoulli-p random subset of P. The value assigned to any lifted monomial $x_u L_{I,J}$ is the corresponding product-measure moment: it is 0 if $u \in J$, $y_{C,I,J}$ if $u \in I$, and $p\, y_{C,I,J}$ if $u \notin I \cup J$. Since these are moments of an actual probability distribution on $\{0,1\}^P$, the lifted box constraints are satisfied; in particular $0 \le y_{C,I,J} \le 1$ for every lifted variable.

Now multiply a covering constraint

$$g_C(x) := \sum_{u \in C} x_u - 1 \ge 0$$

by any literal $L_{I,J}$ with $|I \cup J| \le r$. If $I \cap J \ne \emptyset$, then $L_{I,J} \equiv 0$. Assume $I \cap J = \emptyset$. If $I \cap C \ne \emptyset$, then whenever $L_{I,J} = 1$ at least one variable in C is already fixed to 1, so $g_C(x) L_{I,J}(x) \ge 0$ pointwise.

It remains to handle $I \cap C = \emptyset$. Let $c = |J \cap C|$. Conditional on the event $L_{I,J} = 1$, the c variables in $C \cap J$ are fixed to 0, while the remaining $\ell - c$ variables of C remain independent Bernoulli-p. Hence

$$\mathbb{E}\left[\sum_{u \in C} x_u \,\middle|\, L_{I,J} = 1 \right] = (\ell - c)p.$$

Since $c \le |J| \le |I \cup J| \le r$, we have

$$(\ell - c)p \ge (\ell - r)\frac{1}{\ell - r} = 1.$$

Therefore $\mathbb{E}[g_C(x) L_{I,J}(x)] \ge 0$ for every lifted covering row, so the product Bernoulli moment vector is feasible for the level-r Sherali–Adams relaxation. Its objective value is

$$\mathbb{E} \sum_{u \in P} x_u = Np = \frac{N}{\ell - r}.$$

□

Thus no argument whose entire certificate is bounded by the above level-r Sherali–Adams relaxation can certify a transversal lower bound exceeding $N/(\ell - r)$ on these instances, even though the true integral optimum is $> \alpha N$. The artifact script `scripts/sa_barrier_verification.py` exhaustively checks the lifted covering inequalities for $N \leq 7$ using exact rational arithmetic.

8.2. **A q-shadow dichotomy.** The next obstruction is an extremal statement for Johnson shadows. It explains why a sparse separator family cannot be forced by expansion alone unless a large pre-existing shadow is already present.

Let P be a K-point set, write $K = h + L$ with $L \geq 1$, and fix $1 \leq q \leq h$. Put $X = \binom{P}{q}$ and $Y = \binom{P}{h}$. A q-set Q captures an h-set H if $Q \subseteq H$. Let $R \subseteq Y$ be a live packet of density $r := |R|/|Y| > 0$. Let $D \subseteq X$ be a previously forbidden q-shadow, and let $\mathcal{B} \subseteq Y$ be a family of already played h-sets. Define

$$\delta_q := \frac{\binom{K-q}{h-q}}{\binom{K}{h}},$$

the density of h-sets containing a fixed q-set. Write

$$\partial_h D := \{H \in Y : \text{some } Q \in D \text{ satisfies } Q \subseteq H\}, \qquad \sigma_q(D) := \frac{|\partial_h D|}{|Y|}.$$

Thus $\sigma_q(D)$ is an h-level shadow measure, not a q-level count. The normalized second singular value of the q-versus-h inclusion graph satisfies

$$\lambda_q^2 = \frac{q(K-h)}{h(K-q)}.$$

Proposition 8.2 (q-shadow covering dichotomy)**.** *Let*

$$U_q(D, \mathcal{B}) := D \cup \bigcup_{B \in \mathcal{B}} \binom{B}{q}, \qquad A := X \setminus U_q(D, \mathcal{B})$$

be the excluded and legal q-families. Assume no legal q-separator $Q \in A$ captures at least

$$\frac{1}{2}\,\delta_q |R|$$

members of R. Then

$$\frac{|U_q(D, \mathcal{B})|}{|X|} > 1 - \frac{4\lambda_q^2}{r}.$$

Consequently

$$\sigma_q(D) + |\mathcal{B}|\delta_q > 1 - \frac{4\lambda_q^2}{r}.$$

In particular, if $r \geq \eta > 0$, the right-hand side may be weakened to $1 - 4\lambda_q^2/\eta$.

Proof. Put $a = |A|/|X|$. The case $A = \emptyset$ is immediate, so assume $a > 0$. The expander-mixing inequality for the inclusion graph between X and Y gives

$$\left| e(A, R) - \frac{|A|}{|X|}\frac{|R|}{|Y|} e(X, Y) \right| \leq \lambda_q\, e(X, Y) \sqrt{\frac{|A|}{|X|}\frac{|R|}{|Y|}}.$$

If

$$a \geq \frac{4\lambda_q^2}{r},$$

then the average number of elements of R captured by a member of A is at least half the random expectation, namely at least $\frac{1}{2}\delta_q|R|$. This contradicts the hypothesis that no legal separator captures that many live packet members. Hence

$$a < \frac{4\lambda_q^2}{r},$$

which is exactly

$$\frac{|U_q(D,\mathcal{B})|}{|X|} = 1 - a > 1 - \frac{4\lambda_q^2}{r}.$$

It remains only to convert this exact q-level covering statement into the displayed shadow bound. By the LYM normalized matching property on the Boolean lattice [And87, Section 2.2], every $\mathcal{F} \subseteq \binom{P}{q}$ satisfies

$$\frac{|\partial_h \mathcal{F}|}{\binom{K}{h}} \geq \frac{|\mathcal{F}|}{\binom{K}{q}} \qquad (q \leq h).$$

Applying this to $\mathcal{F} = D$ gives $|D|/|X| \leq \sigma_q(D)$. Also each $B \in \mathcal{B}$ contains exactly $\binom{h}{q}$ members of X, so $\binom{h}{q}/\binom{K}{q} = \delta_q$. The union bound gives

$$\frac{|U_q(D,\mathcal{B})|}{|X|} \leq \frac{|D|}{|X|} + |\mathcal{B}|\delta_q \leq \sigma_q(D) + |\mathcal{B}|\delta_q,$$

and the conclusion follows. □

In the central regime relevant to top-facet packet arguments, take $L \sim h/\log h$ and $q \sim 2(\log h)^2$. Then

$$\delta_q = (e + o(1))h^{-2}, \qquad \lambda_q^2 = (2 + o(1))\frac{\log h}{h}.$$

Thus, when r is bounded below and $\sigma_q(D) = o(1)$, separator starvation forces $|\mathcal{B}| > (e^{-1} - o(1))h^2$. Pure expansion cannot supply an unlimited sequence of fresh separators.

8.3. **A separator-only limitation.** The last obstruction is a finite model for strategies which only play odd carrier separators and ignore dyadic shielding.

Proposition 8.3 (Separator-only limitation)**.** *Let*

$$U_n = (n/2, n] \cap \mathbb{Z}, \qquad r_1(n) = \frac{n(\log\log n)^2}{\log n}.$$

For each sufficiently large n, let $S_n \subseteq [2, n/2]$ be a finite set of odd integers such that

$$|S_n| = o(r_1(n)) \quad \textit{and} \quad H(S_n) := \sum_{p \in P(S_n)} \frac{1}{p} = o(1),$$

where $P(S_n)$ is the set of odd primes dividing at least one element of S_n. Consider any separator-only closure in which the only lower-half elements allowed to kill members of U_n are Shortener-claimed elements of S_n; upper-half elements may only be removed individually after they have actually been played. Then Prolonger has a legal play prefix of length $o(r_1(n))$ after which every separator in S_n is dead, while at least

$$\frac{n}{2} - o(n)$$

elements of U_n remain legal. Therefore this finite odd-carrier separator-only proof class cannot prove $L(n) = O(r_1(n))$, or indeed any $o(n)$ upper bound.

Proof. Fix n and write $S = S_n$ and $P = P(S_n)$. Suppose the current chosen set is an antichain A, every lower-half Shortener claim lies in S, and some $s \in S$ is still legal. Since $s \le n/2$, choose $a \ge 0$ maximal with $2^a s \le n$, and put

$$x = 2^a s.$$

Then $x \in U_n$, because maximality gives $2^{a+1}s > n$ and hence $2^a s > n/2$.

We claim that x is legal. First, any previously chosen upper-half element $y \in A \cap U_n$ is incomparable with x unless $x = y$, since two distinct integers larger than $n/2$ cannot divide one another. The equality $x = y$ is also impossible: if x had already been chosen, then $s \mid y$, contradicting the assumption that s is still legal. Second, let $t \in A \setminus U_n$. By the definition of the proof class, $t \in S$ and is odd. Since $x > n/2 \ge t$, we cannot have $x \mid t$. If $t \mid x = 2^a s$, then oddness of t implies $t \mid s$, so t is comparable with the still-legal separator s, again impossible. Thus x is a legal Prolonger move.

After Prolonger plays x, the separator s is dead forever because $s \mid x$. Repeating this whenever a legal element of S remains, Prolonger kills at least one still-legal separator on each preemption move, while Shortener's separator-only moves can only kill additional members of S. After at most $|S|$ Prolonger preemptions and at most $|S|$ intervening Shortener moves, no element of S remains legal. The prefix length is therefore at most $2|S| = o(r_1(n))$.

It remains to count legal upper-half integers after this prefix. Let $A_U = A \cap U_n$. The prefix has length $O(|S|)$, so $|A_U| = o(n)$. An unplayed upper-half integer $u \in U_n \setminus A_U$ can be illegal because of a lower-half separator only if some selected $t \in S$ divides u. In that case u is divisible by at least one prime $p \in P$. Hence the number of upper-half integers killed by lower-half separators is at most

$$\left|\bigcup_{p\in P}\{u \in U_n : p \mid u\}\right| \le \sum_{p\in P}\left(\frac{n}{p}+1\right) \le 2n\sum_{p\in P}\frac{1}{p} = 2nH(S) = o(n),$$

where we used $p \le n$, so $1 \le n/p$. Therefore at least

$$|U_n| - |A_U| - o(n) = \frac{n}{2} - o(n)$$

upper-half integers remain legal. Since $r_1(n) = o(n)$, the separator-only closure cannot certify an $O(r_1(n))$, or more generally an $o(n)$, upper bound. □

The packet separator families motivating this proposition are a special case. Let $h = \lfloor \log\log n \rfloor$, $Y_0 = n^{1/(2h)}$, let $P_n \subseteq [Y_0, 2Y_0]$ be a prime packet, and fix $q_0 \le h/4$. The carrier family

$$S_n = \left\{\prod_{p\in Q} p :\ Q \subseteq P_n,\ 1 \le |Q| \le q_0\right\}$$

has the support-covering relation

$$\prod_{p\in Q} p \mid 2^a \prod_{p\in T} p \quad\Longleftrightarrow\quad Q \subseteq T.$$

For large n, every carrier in S_n is odd and at most $(2Y_0)^{q_0} = n^{1/8+o(1)} < n/2$; moreover $|S_n| \le (q_0+1)(2Y_0)^{q_0} = n^{1/8+o(1)} = o(r_1(n))$, and the prime number theorem gives

$$H(S_n) \le \sum_{p\in P_n}\frac{1}{p} \le \frac{|P_n|}{Y_0} = O\left(\frac{1}{\log Y_0}\right) = O\left(\frac{\log\log n}{\log n}\right) = o(1).$$

Thus the theorem applies to the finite packet closure. Its scope is deliberately class-level: it does not rule out separator methods which also track dyadic lifts, exact upper targets, or composite fallback moves outside the finite odd-carrier layer.

9. Conclusion

The main unconditional bounds of Theorems 4.4 and 7.21 bracket the divisibility antichain saturation game as

$$\left(\frac{1}{8}-o(1)\right)\frac{n\log\log n}{\log n}\le L(n)\le\left(\frac{\mathcal{W}_4}{2}+o(1)\right)n.$$

The lower bound comes from a two-phase fan-capture strategy, while the upper bound improves the elementary linear picture by using a lower-profile envelope for Shortener's legal-prime prefix and a fourth-order Bonferroni comparison.

The conditional T2 theorem points toward a substantially larger lower-bound scale, but it is deliberately stated under a restricted safe-edge hypothesis: the general safe-edge property is false even inside the residual arithmetic games. The finite graph, hypergraph, and embedding components are robust; the remaining issue is whether the safe-edge property holds for states reached under the Prolonger activation strategy.

On the structural side, the shield reduction and Theorem A show that short upper-half shield prefixes cannot prove a linear game lower bound through the weighted shadow functional $\beta_n(P)$.

Three obstructions in Section 8 explain why the next improvements are unlikely to come from purely local relaxations, fixed shadow-cover dichotomies, or separator-only strategies. A sharper upper bound will need to use dynamic information about the actual divisibility game, not merely the static incidence geometry of low-rank shadows.

The central question remains the sharp order of growth of $L(n)$, and in particular whether $L(n)=o(n)$. The methods here give new constants, exact auxiliary structure, and a set of formal targets for future work, but they do not settle that asymptotic dichotomy.

Appendix A. Deferred Proofs

A.1. **Conditional proof of the T2 lower bound.** This appendix proves the conditional lower bound in Theorem 4.7. The proof is self-contained: all target families, auxiliary games, weights, activation rules, deletion errors, and asymptotic estimates are defined below. The opening counterexamples explain why the restricted finite safe-edge hypothesis is stated explicitly rather than hidden inside a max-gain argument.

Throughout this appendix n is large, $0<\delta<1/4$ is fixed, and

$$Y:=n^{\delta}.$$

Let $\mathcal{P}_Y$ be the set of odd primes at most Y. For $a<c$ in $\mathcal{P}_Y$, put

$$I_{a,c}:=\left(\frac{n}{2ac},\frac{n}{ac}\right],\qquad B_{a,c}:=I_{a,c}\cap\mathbb{P},$$

and define the target family

$$\mathcal{T}:=\{\,acb:\ a<c,\ a,c\in\mathcal{P}_Y,\ b\in B_{a,c}\,\}.$$

For large n, every such b is greater than Y: indeed $b>n/(2Y^2)=\frac{1}{2}n^{1-2\delta}>Y$. Thus a,c,b are distinct primes and each target acb lies in $(n/2,n]$.

Auxiliary slot game and potential. The scored rank-three slot hypergraph $G(n,\delta)$ has vertex set consisting of the slots b,ab,cb which occur in the targets above, and has weighted edges

$$e_{a,c,b}:=\{b,ab,cb\},\qquad a<c,\quad a,c\in\mathcal{P}_Y,\quad b\in B_{a,c},$$

each of weight 1. More generally, the finite auxiliary games below are weighted incidence games of rank at most 3. A state has a captured-vertex set C, a deleted-vertex set D_V, an unscored deleted-edge set D_E, a scored edge set K, and total scored weight S. An edge is live if it is not in $D_E\cup K$ and contains no deleted vertex. Maker may capture a live edge, adding its vertices to C and adding its weight to S; in the scored hypergraph variant, Maker may instead use an *alternate-scoring move* to score the target without adding slot vertices to C. Breaker may reply by deleting one uncaptured vertex, by deleting one live edge without scoring it, or in the scored-hypergraph variant by making a scored-edge reply, which adds one live edge to K and increments S by its weight. The scaled potential is

$$Q := 8S + \sum_{e \text{ live}} 2^{|e\cap C|} w(e).$$

Proposition A.1 (Max-gain domination is false in rank three). *The following finite scored rank-3 capture state has a unique max-gain Maker move, but Breaker can answer that move by deleting one uncaptured vertex and make the scaled potential decrease.*

Proof. Let the already-captured vertices be

$$C = \{z_{ij} : 1 \le i, j \le 3\},$$

and let u_1, u_2, u_3, v be uncaptured. There are no deleted or scored edges. The live edges are

$$f = \{u_1, u_2, u_3\}, \qquad w(f) = 9,$$

and, for $1 \le i, j \le 3$,

$$e_{ij} = \{v, u_i, z_{ij}\}, \qquad w(e_{ij}) = 4.$$

For a live edge e, write $\Phi(e) = 2^{|e\cap C|} w(e)$ and $Q = 8S + \sum \Phi(e)$.

Before Maker moves, $\Phi(f) = 9$ and $\Phi(e_{ij}) = 8$ for all i, j. If Maker plays f, then the score contribution of f changes by $8w(f) - \Phi(f) = 72 - 9 = 63$, and each of the nine edges e_{ij} doubles from 8 to 16. Hence

$$\Delta(f) = 63 + 9 \cdot 8 = 135.$$

After this move, the vertex v is still uncaptured, and deleting v removes all nine edges e_{ij}, now of potential 16 each. The Breaker loss is

$$9 \cdot 16 = 144 > 135.$$

Thus the full round decreases Q.

It remains only to check that f is the unique max-gain move. If Maker plays some e_{ij}, its own score gain is $8\cdot 4 - 8 = 24$. The edge f doubles from 9 to 18, giving gain 9. The two other edges with the same u_i gain 24 each, and the six edges with the other two u-vertices gain 8 each. Therefore

$$\Delta(e_{ij}) = 24 + 9 + 2 \cdot 24 + 6 \cdot 8 = 129 < 135.$$

So the unique max-gain move is f, and it does not dominate the subsequent vertex deletion at v. ☐

Proposition A.2 (K_4-fiber refutation in the residual arithmetic game). *For every fixed $0 < \delta < 1/4$ and all sufficiently large n, the residual slot hypergraph $G(n, \delta)$ contains a reachable pre-Maker state at which every legal Maker move admits a legal Breaker reply that strictly decreases the scaled potential Q. Hence the general safe-edge hypothesis is false even inside the arithmetic games $G(n, \delta)$.*

Proof. Here “reachable” means reachable under some legal play sequence, not necessarily under the specific Prolonger strategies of Proposition A.9. The construction refutes the universal safe-edge hypothesis over the full state space; the main theorem’s restricted hypothesis Definition 4.5 is formulated for states generated by Prolonger’s specific strategies, not over the full state space.

Choose the four fixed small primes $13, 17, 19, 23$. Their pairwise products are

$$13 \cdot 17 = 221, \quad 13 \cdot 19 = 247, \quad 13 \cdot 23 = 299, \quad 17 \cdot 19 = 323, \quad 17 \cdot 23 = 391, \quad 19 \cdot 23 = 437.$$

All six products lie between 221 and 437, and $2 \cdot 221 = 442 > 437$. For sufficiently large n, since $23 \le n^\delta = Y$, all four primes lie in $\mathcal{P}_Y$. By the prime number theorem, the interval

$$\left(\frac{n}{442}, \frac{n}{437}\right]$$

contains a prime q for all sufficiently large n. Also $q > Y$ for large n, since $\delta < 1/4$.

For every pair $a < c$ among $\{13, 17, 19, 23\}$, we have

$$q > \frac{n}{442} \geq \frac{n}{2ac} \qquad \text{and} \qquad q \leq \frac{n}{437} \leq \frac{n}{ac},$$

because $ac \geq 221$ and $ac \leq 437$. Hence $q \in B_{a,c}$ for all six pairs, and the six edges

$$e_{a,c} := \{q, aq, cq\}$$

all occur in $G(n, \delta)$. They form a K_4-fiber over the common slot q.

We next reach a state in which only these six edges are live and no vertex in them has yet been captured. While any other edge is live, Maker uses the alternate-scoring move to score that exact target without adding slot vertices to C. Any edge outside this six-edge family has at least one slot outside

$$U_{\text{fib}} := \{q, 13q, 17q, 19q, 23q\},$$

so Breaker deletes such an outside slot. This never harms the six desired edges. Repeating this finite operation leaves exactly the six K_4-fiber edges live, with $C = \varnothing$.

From that state, play two legal rounds. First Maker captures

$$e_{13,17} = \{q, 13q, 17q\},$$

so $q, 13q, 17q \in C$, and Breaker deletes the live edge $e_{13,19}$. Next Maker captures

$$e_{19,23} = \{q, 19q, 23q\},$$

so $q, 13q, 17q, 19q, 23q \in C$, and Breaker deletes the live edge $e_{13,23}$.

At the next pre-Maker state, the only remaining live edges from this fiber are

$$e_{17,19} = \{q, 17q, 19q\}, \qquad e_{17,23} = \{q, 17q, 23q\}.$$

Both have all three vertices in C, so $\Phi(e) = 2^3 \cdot 1 = 8$ for each. There are no other live edges by construction.

Maker has only two possible live edges to play. If Maker plays $e_{17,19}$, by ordinary capture or by alternate scoring, then the change in Q is

$$+8 \cdot 1 - 8 = 0,$$

because S increases by 1, the edge of potential 8 is removed, and no new vertex is added to C. The remaining live edge $e_{17,23}$ still contributes 8. Breaker then deletes $e_{17,23}$, decreasing Q by 8. Thus $Q_{\text{after}} = Q_{\text{before}} - 8$. The same argument holds if Maker instead plays $e_{17,23}$.

Therefore this reachable pre-Maker state has the claimed property: every legal Maker move admits a legal Breaker reply that strictly decreases Q. Note that Breaker's final reply here is an unscored edge deletion, adding $e_{17,23}$ to D_E, not a scored-edge play; under the scored-edge rule, the same move would change Q by 0, not by -8. The refutation shows that the general safe-edge hypothesis fails even when Breaker has access to the unscored-edge-deletion move, which is the regime relevant to the rank-3 residual hypergraph game. ☐

Proposition A.3 (Conditional finite scaled capture)**.** *Consider a finite weighted incidence game of rank at most* 3*, equipped with the scaled potential Q and the Maker–Breaker rules described above. In each full Maker–Breaker round, Breaker may play any one of:*

- (i) *a vertex deletion: add an uncaptured vertex $v \in V \setminus C$ to D_V;*
- (ii) *an unscored edge deletion: add one currently live edge e to D_E, available in both the graph and scored-hypergraph variants;*
- (iii) *a scored-edge reply: add one currently live edge e to K and increment S by $w(e)$, available only in the scored-hypergraph variant.*

The general safe-edge hypothesis stated below is provably false even for $G(n, \delta)$; see Proposition A.2. In this paper the finite-capture proposition is used only under the restricted hypothesis that the same property holds for states actually generated by the activation and residual strategies in Propositions A.6 and A.9. This restricted hypothesis remains open.

Assume the following restricted safe-edge hypothesis: at every pre-Maker state generated by the strategy under consideration with at least one positive-weight live edge, Maker has a legal live edge f such that every legal Breaker reply from (i)–(iii) applicable to the current variant leaves Q at least as large as it was before Maker's move.

Under this restricted hypothesis, if the game is continued until no positive-weight live edge remains, the final claimed or scored weight is at least one eighth of the initial total edge-weight:

$$S_{\mathrm{fin}} \geq \frac{1}{8}\sum_e w(e).$$

Proof. Maker always chooses an edge supplied by the restricted safe-edge hypothesis. By definition of that hypothesis, Q is nondecreasing after every full Maker–Breaker round.

Initially $C = \varnothing$, $S = 0$, and all initial edges are live, so

$$Q_0 = \sum_e w(e).$$

At terminal time there is no positive-weight live edge, hence $Q_{\mathrm{fin}} = 8S_{\mathrm{fin}}$. Therefore

$$8S_{\mathrm{fin}} = Q_{\mathrm{fin}} \geq Q_0 = \sum_e w(e),$$

which is the claimed one-eighth bound.

For the graph variant, a Breaker edge-deletion reply is included in the restricted safe-edge hypothesis. For the scored hypergraph variant, a Breaker scored-edge reply is harmless whenever it occurs, since it changes Q by $8w(e) - 2^{|e\cap C|}w(e) \geq 0$ for rank at most 3; the only nontrivial part of the hypothesis is therefore the post-Maker vertex-deletion bound. □

Proposition A.4 (Initial target mass)**.** *For every fixed* $0 < \delta < 1/4$*,*

$$W_0 := \sum_{\substack{a<c\\ a,c\in\mathcal{P}_Y}} |B_{a,c}| \gg_\delta \frac{n(\log\log n)^2}{\log n}.$$

Proof. Since $ac \leq Y^2 = n^{2\delta}$, the interval endpoint $X_{a,c} := n/(ac)$ satisfies $X_{a,c} \geq n^{1-2\delta} \to \infty$. By the prime number theorem, uniformly for all such pairs and all sufficiently large n,

$$\pi(X_{a,c}) - \pi(X_{a,c}/2) \geq \kappa_\delta \frac{X_{a,c}}{\log n} = \kappa_\delta \frac{n}{ac\log n}$$

for some positive constant κ_δ. Therefore

$$W_0 \geq \kappa_\delta \frac{n}{\log n} \sum_{\substack{a<c\\ a,c\in\mathcal{P}_Y}} \frac{1}{ac}.$$

Mertens' theorem for primes gives

$$\sum_{p\in\mathcal{P}_Y} \frac{1}{p} = \log\log Y + O(1) = \log\log n + O_\delta(1),$$

and $\sum_p p^{-2} < \infty$. Hence

$$\sum_{a<c\in\mathcal{P}_Y} \frac{1}{ac} = \frac{1}{2}\left[\left(\sum_{p\in\mathcal{P}_Y}\frac{1}{p}\right)^2 - \sum_{p\in\mathcal{P}_Y}\frac{1}{p^2}\right] = \left(\frac{1}{2} + o(1)\right)(\log\log n)^2.$$

Combining the last two displays proves the claim. □

Proposition A.5 (Residual divisibility embedding)**.** *Let $\mathcal{T}_*$ be a subfamily of targets $t = acb \in \mathcal{T}$ with the following secured-pair origin. For every $t = acb \in \mathcal{T}_*$, there is a securing target*

$$t_0(a, c) = acb_0 \in \mathcal{T}, \qquad b_0 \neq b,$$

which has already been played, making a, c, and ac unavailable as move options; and no other earlier move is comparable with any $t \in \mathcal{T}_*$. Attach to t the three slot labels

$$b, \qquad ab, \qquad cb$$

and let $e_t := \{b, ab, cb\}$. Then the residual divisibility game on $\mathcal{T}_*$ is at least as favorable to Prolonger as the scored rank-3 hypergraph game on the hyperedges $\{e_t : t \in \mathcal{T}_*\}$:

1. *the only future moves that can make $t = acb$ illegal are b, ab, cb, and t itself;*
2. *playing b, ab, or cb deletes exactly the residual targets whose hyperedges contain that slot;*
3. *playing the exact target t kills no other target in $\mathcal{T}_*$, and is therefore a scored-edge move;*
4. *every live hyperedge corresponds to a legal actual move t.*

Proof. Fix $t = acb \in \mathcal{T}_*$. Because $t > n/2$, no proper multiple of t lies in $\{2, \dots, n\}$. Thus a future move can make t illegal only by being t itself or a proper divisor of t. The positive divisors are

$$1,\ a,\ c,\ ac,\ b,\ ab,\ cb,\ acb.$$

The divisor 1 is not a move, while a, c, and ac are unavailable by the played securing target $t_0(a, c)$. This proves the first assertion.

For slot incidence, suppose another target is $t' = a'c'b'$. Since $b, b' > Y \geq a, a', c, c'$, the divisibility $b \mid t'$ holds iff $b = b'$. Thus a move b kills exactly the targets in the b-fiber. Similarly, if $ab \mid a'c'b'$, the large prime factor forces $b = b'$, and after cancellation a must be a' or c'. This is exactly the condition that the slot ab belongs to $e_{t'}$. The proof for cb is identical.

If two targets t, u in $\mathcal{T}_*$ are comparable, say $t \mid u$, then u is a multiple of t lying in $(n/2, n]$. The next positive multiple of t after t is $2t > n$, so $u = t$. Hence distinct targets are pairwise incomparable, and playing an exact target scores only that target.

Finally, if e_t is live in the slot hypergraph, then none of b, ab, cb, or t has been played. The only other proper divisors of t are $1, a, c, ac$, and these are either not moves or are unavailable because of the played securing target $t_0(a, c)$. The securing target is a distinct target larger than $n/2$, hence is incomparable with t, and there is no proper multiple of t on the board. By hypothesis, no other earlier move is comparable with any target in $\mathcal{T}_*$. Therefore no played number is comparable with t, so the actual move t is legal. This proves the comparison. ∎

Proposition A.6 (Activation-stage potential)**.** *During activation, let each unclaimed pair edge $e = (a, c)$ carry the current token set $B_e(t)$ of primes $b \in B_{a,c}$ for which the target acb has not yet been killed. Write $w_t(e) := |B_e(t)|$.*

Assume the restricted safe-edge hypothesis of Proposition A.3 for the activation graph states generated below. Prolonger plays as follows. In each activation round, among pair edges with positive current token weight, he chooses a safe edge supplied by that hypothesis; if $e = (a, c)$ is chosen, he plays one live target acb with $b \in B_e(t)$. This scores one actual move, claims the pair edge e, captures the vertices a and c, and leaves the remaining live tokens on that secured pair for the residual phase.

Let S_t be the number of activation targets already played by Prolonger, and define

$$Q_t := S_t + \sum_e \phi_t(e),$$

where

$$\phi_t(e) = \begin{cases} \frac{1}{8}w_t(e), & e \textit{ unclaimed and neither endpoint captured,} \\ \frac{1}{4}w_t(e), & e \textit{ unclaimed and one endpoint captured,} \\ \frac{1}{2}w_t(e), & e \textit{ unclaimed and two endpoints captured,} \\ w_t(e), & e \textit{ claimed,} \\ 0, & e \textit{ deleted.} \end{cases}$$

Let E be the total number of target tokens deleted during activation by Shortener moves not modeled as graph vertex deletions a or graph edge deletions ac. Here E counts the total number of live target tokens deleted by off-model Shortener moves during activation, not the number of such moves: a single off-model move can delete multiple tokens, and E tallies deletions, not move events. At the end of activation,

$$Q_{\text{end}} \geq \frac{W_0}{8} - E.$$

Equivalently,

$$S_{\text{end}} + \sum_{e \text{ secured}} w_{\text{end}}(e) \geq \frac{W_0}{8} - E.$$

Proof. At each activation state, form the finite rank-2 graph game whose vertices are the small primes, whose live edges are the unclaimed pairs $e = (a, c)$ with current weight $w_t(e) > 0$, and whose claimed, deleted, and captured objects are the secured pairs, exhausted pairs, and small primes already made unavailable. The restricted safe-edge hypothesis is applied to this dynamically updated weight function $w_t(e)$ at the current pre-Prolonger state. A modeled Shortener reply of type a or c is exactly a graph vertex deletion, and a modeled reply of type ac is exactly a graph edge deletion.

First ignore off-model token deletions. Under the identification above, the activation process is the rank-2 graph game of Proposition A.3, with one bookkeeping change: when Prolonger claims an edge e, he consumes one token from e and adds it immediately to the score S_t. If the coefficient of e before the move is $\gamma \in \{1/8, 1/4, 1/2\}$, the contribution of e changes from $\gamma w_t(e)$ to

$$1 + (w_t(e) - 1) = w_t(e),$$

exactly as if the edge had become claimed with full weight. All other incident edges change as in the graph game. Thus the scaled potential argument gives $Q_{\text{end}} \geq W_0/8$ when every harmful Shortener move is a graph deletion.

Now restore all actual Shortener moves. A token acb is deleted precisely when Shortener plays one of its still-available harmful divisors or the target itself. Deleting one live token can reduce Q_t by at most 1, since every coefficient appearing in the definition of $\phi_t(e)$ is at most 1. Therefore all off-model deletions together reduce the lower bound by at most E, giving

$$Q_{\text{end}} \geq W_0/8 - E.$$

At the end of activation, a positive-weight edge is either secured or deleted, so Q_{end} is exactly the activation score plus the residual live target weight on secured pairs. □

Proposition A.7 (Activation deletion budget)**.** *For fixed $0 < \delta < 1/4$, the off-model deletion count from Proposition A.6 satisfies*

$$E \ll \frac{Y^4}{\log^4 Y} = o\left(\frac{n(\log\log n)^2}{\log n}\right).$$

Proof. There is at most one activation round per small-prime pair, so the number of rounds is

$$R \leq \binom{|\mathcal{P}_Y|}{2} \ll \frac{Y^2}{\log^2 Y}.$$

We bound the number of live tokens an off-model Shortener move can delete.

If Shortener plays a large prime b, then for each pair (a, c) there is at most one token acb using that b. Hence the move deletes at most $O(Y^2/\log^2 Y)$ tokens. If Shortener plays a lateral semiprime pb with $p \le Y$ and $b > Y$ prime, the killed targets have the form pcb with the third small prime varying, so there are at most $O(Y/\log Y)$ of them. If Shortener plays an exact target acb, exactly one token is deleted.

The worst of these bounds is the large-prime bound. Across R activation rounds,

$$E \ll R\frac{Y^2}{\log^2 Y} \ll \frac{Y^4}{\log^4 Y}.$$

Since $Y = n^\delta$,

$$\frac{Y^4/\log^4 Y}{n(\log\log n)^2/\log n} \ll_\delta \frac{n^{4\delta-1}}{(\log n)^3(\log\log n)^2} \to 0$$

because $4\delta - 1 < 0$. □

Proposition A.8 (Residual mass on secured pairs)**.** *At the end of activation, the secured pairs carry residual live target weight*

$$M := \sum_{e \text{ secured}} w_{\text{end}}(e) \gg_\delta \frac{n(\log\log n)^2}{\log n}.$$

Proof. By Proposition A.6,

$$S_{\text{end}} + M \ge \frac{W_0}{8} - E.$$

The activation score is at most the number of small-prime pairs:

$$S_{\text{end}} \le R \ll \frac{Y^2}{\log^2 Y} = o\left(\frac{n(\log\log n)^2}{\log n}\right),$$

because $2\delta < 1$. The deletion budget E is negligible by Proposition A.7, while W_0 has the required lower bound by Proposition A.4. Subtracting the two negligible terms gives the claim. □

Proposition A.9 (Conditional T2 lower bound)**.** *Assume the restricted safe-edge hypothesis of Proposition A.3 for the activation graph states and for the residual slot-hypergraph states generated by the construction above. For every fixed $0 < \delta < 1/4$ there is $c_\delta > 0$ such that*

$$L(n) \ge c_\delta \frac{n(\log\log n)^2}{\log n}$$

for all sufficiently large n.

Proof. After activation, restrict to the residual live targets lying over secured pairs. Their total weight is M, as in Proposition A.8. For each secured pair (a, c), the activation move supplies a securing target $t_0(a, c) = acb_0$, and the residual subfamily keeps only targets acb with $b \ne b_0$ which remain live. Thus the securing target makes a, c, and ac unavailable while being incomparable with every remaining residual target over the same pair. Modeled activation replies produce integers that may be comparable with some residual targets in $\mathcal{T}_*$; any such comparable target has already been removed from $\mathcal{T}_*$ by the residual construction. Hence no live hyperedge in the rank-3 hypergraph corresponds to a target killed by a modeled reply. Any off-model move comparable with a target deletes that target and so removes it from the residual live subfamily. Therefore the strengthened hypothesis of Proposition A.5 compares the remaining divisibility game to a scored rank-3 slot hypergraph whose total edge weight is M.

Under the residual restricted safe-edge hypothesis, Proposition A.3 gives a strategy that forces scored residual target weight at least $M/8$ in this hypergraph model. The comparison proposition says that every such scored hypergraph move is a genuine actual game move, and that the actual game gives Prolonger no fewer legal options than the model. Hence the number of residual moves after activation is at least $M/8$. Since $M \gg_\delta n(\log\log n)^2/\log n$, the same lower bound holds for the total game length, conditional on the restricted safe-edge hypothesis.

This proves the per-δ statement Theorem 4.7; the intro's absolute-c corollary follows by specializing $\delta = 1/8$. □

A.2. **Diagonal selection conditions.** This subsection records the eventual side conditions used in Proposition 7.19. In the notation of that proposition, after $H_m, \lambda_m, a_m, \eta_m, s_m$ have been fixed, choose N_m recursively increasing so that for every $n \geq N_m$ the following twelve conditions hold:

(1) the prefix-existence conclusion of Lemma 7.1 holds up to index K;
(2) the local-density error from Proposition 7.3 satisfies

$$\xi_{H_m}(n) \leq \frac{1}{m}; \tag{7.10a}$$

(3) the envelope moment estimate from Proposition 7.8 satisfies

$$|T_r^{(b(H_m))}(n) - J_r^{(H_m)}| \leq \frac{1}{4m} \qquad (1 \leq r \leq 4); \tag{7.10b}$$

(4) the bridge moment estimate from Proposition 7.18 satisfies

$$|T_r^{(p(H_m))}(n) - T_r^{(b(H_m))}(n)| \leq \frac{1}{2m} \qquad (1 \leq r \leq 4); \tag{7.10c}$$

(5) the cell-reciprocal-mass estimate of Lemma 7.11 holds for the finitely many fixed-ratio cells used with parameters H_m, a_m, s_m, with all limiting errors bounded by $1/m$;
(6) the product-strip estimate of Lemma 7.12 holds for $C = \lambda_m^r$ and $1 \leq r \leq 4$, with associated finite cell-count estimates valid to tolerance $1/m$;
(7) the queue-clearing Hall inequalities in Proposition 7.18, with $a = a_m$, $\eta = \eta_m$, and $s = s_m$, hold for every suffix $0 \leq t \leq s_m$;
(8) the exceptional-index reciprocal mass estimate in (7.6) holds in the form

$$\sum_{j\in E} \frac{1}{b_j} \leq \frac{1}{m};$$

(9) all fixed-ratio prime number theorem estimates used in the queue construction, in particular (7.4), hold with tolerance $1/m$ for the finitely many ratios determined by a_m and s_m;
(10) the smallest genuine scale satisfies

$$n^{\alpha_{H_m-1}} = n^{1/H_m+\tau_{H_m}} \geq m; \tag{7.10d}$$

(11) the interblock nonoverlap inequalities

$$a_m^2 B_h < A_{h-1} \qquad (2 \leq h \leq H_m - 1), \qquad a_m^2 B_1 < n$$

hold for the block endpoints associated to H_m;
(12) the top-block prime supply satisfies

$$\pi(2n) - \pi(n) \geq K.$$

Each condition is eventual because $H_m, \lambda_m, a_m, \eta_m, s_m$ and the finite list of ratios are fixed while $n \to \infty$.

A.3. **Restricted carrier classes.** The next two propositions record unconditional upper bounds for two restricted families of Prolonger carrier moves. They are not used in the main upper bound, but they clarify which carrier configurations are not responsible for the remaining difficulty.

Proposition A.10 (Small-prime composite Prolonger with disjoint supports)**.** *Fix $1/3 < \alpha < 1/2$ and put $y = n^{\alpha}$. Suppose that every composite Prolonger move has all prime factors at most y, and that the prime supports of distinct composite Prolonger moves are pairwise disjoint. Against this restricted class, Shortener has a strategy forcing*

$$L(n) = O_{\alpha}\left(\frac{n}{\log n}\right).$$

Proof. Shortener uses the following priority list, always skipping moves which have already become illegal. First, she plays every legal prime. After this prime phase, define B to be the set of primes $p \le y$ whose singleton was not resolved by prime play because p appears in an observed composite Prolonger carrier. By the disjoint-support hypothesis, each $p \in B$ belongs to a unique composite carrier $C(p)$.

Second, for every $p \in B$ whose carrier has at least two distinct prime factors, write $e(p) = v_p(C(p))$. If $p^{e(p)+1} \le n$, Shortener plays $p^{e(p)+1}$. This move is legal when it is attempted: it is not comparable with $C(p)$, because it has a larger p-adic exponent while $C(p)$ contains another prime factor; it is not comparable with any other composite carrier by disjointness of supports; and previous priority moves either contain different prime supports or have smaller rank.

Third, for every pair $p, q \in B$ belonging to different carriers, Shortener plays the legal squarefree semiprime pq when possible. Such a move is incomparable with every carrier, since no carrier contains both p and q, and neither p nor q alone is available as a legal prime at this stage.

We use the following skipped-move invariant. Under the disjoint-support hypothesis, no prior composite Prolonger move is a proper multiple of a planned Shortener move in the second or third phase. For the phase-two move $p^{e(p)+1}$, the unique carrier containing p has only p-adic exponent $e(p)$, while all other carriers omit p. For the phase-three move pq, the two primes lie in different carriers, so no single carrier contains both. Thus when a planned move is skipped because it is already illegal, the earlier comparable move may be treated as a divisor of the planned move, or the planned move itself, but not as a proper composite multiple coming from Prolonger.

The number of moves in the three phases is

$$O(\pi(n)) + \pi(y) + \pi(y)^2 = O_{\alpha}\left(\frac{n}{\log n}\right),$$

because $2\alpha < 1$.

We now show that after these phases every future legal move has no prime factor at most y. Let x be a legal move divisible by a prime $p \le y$.

If $p \notin B$, then phase one either played p or p was already played by Prolonger as a prime; in either case any multiple of p is illegal. Thus $p \in B$.

If x contains primes from two distinct carriers, then one of the phase-three semiprimes pq divides x, unless it was already illegal because an earlier played move dividing pq also divides x. Either way x is illegal.

Hence all prime factors of x which are at most y lie in a single carrier support. If the carrier is a pure prime power p^e, then any x divisible by p and supported only on this carrier is another power of p, and is comparable with p^e. If the carrier has at least two distinct prime factors, then either x divides the carrier, in which case x is illegal, or for some p it has exponent exceeding $v_p(C(p))$. In the latter case $p^{e(p)+1} \mid x$, so the phase-two move makes x illegal.

Thus all primes $\le y$ are resolved. Any remaining legal integer has all prime factors greater than y. Since $y^3 > n$, it has at most two prime factors. The number of integers at most n with one or two prime factors, all exceeding $y = n^{\alpha}$, is $O_{\alpha}(n/\log n)$ by the prime number theorem and the standard semiprime estimate

$$\sum_{p>y} \pi(n/p) \ll_{\alpha} \frac{n}{\log n} \sum_{p>y} \frac{1}{p} = O_{\alpha}\left(\frac{n}{\log n}\right),$$

where the sum is over $p \le n/y$. Adding the phase moves proves the claim. □

Proposition A.11 (Squarefree rank-three small-prime carriers)**.** *Fix $1/3 < \alpha < 1/2$ and put $y = n^\alpha$. Suppose that every composite Prolonger move is squarefree, has all prime factors at most y, and has support size at most 3. Against this restricted class, Shortener has a strategy forcing*

$$L(n) = O_\alpha\left(\frac{n}{\log n}\right).$$

Proof. Let B be the set of primes which occur in Prolonger's composite carriers. Shortener again uses a fixed priority list, skipping illegal moves.

First she plays every legal prime. Second she plays p^2 for every $p \in B$ when legal. Third she plays every legal squarefree semiprime pq with $p, q \in B$. Fourth she plays every legal squarefree triple pqr with $p, q, r \in B$.

The skipped-move invariant is as follows. Under the squarefree rank-≤ 3 restriction, a prior Prolonger move which makes a planned square or planned triple illegal cannot be a proper multiple of that planned move: a square has no squarefree proper multiple, and a proper multiple of a squarefree triple has rank at least 4. For a planned semiprime pq, the only possible proper multiple is a rank-three carrier pqr; in the support-size-2 case below the remaining candidate is exactly pq, hence it is already illegal because $pq \mid pqr$. Thus the later skipped-move arguments may charge a skipped triple to a prior divisor or to the triple itself, and skipped semiprimes are harmless in the only place where they are used.

The first phase has $O(n/\log n)$ moves, the second has at most $\pi(y)$ moves, and the third has at most $\pi(y)^2 = o(n/\log n)$ moves. It remains to bound the number of triples

$$T_\alpha(n) := \#\{p < q < r \le y : \ pqr \le n\}.$$

Split according to whether $pq \le n^{1-\alpha}$.

If $pq \le n^{1-\alpha}$, then there are $O(n^{1-\alpha} \log\log n/\log n)$ possible pairs $p < q$, and for each pair at most $\pi(y) = O(n^\alpha/\log n)$ choices of r. This contributes $o(n/\log n)$.

If $pq > n^{1-\alpha}$, then $q \le y = n^\alpha$ implies $p > n^{1-2\alpha}$. In this range,

$$\#\{r : \ q < r \le y, \ pqr \le n\} \ll_\alpha \frac{n}{pq \log n}.$$

Therefore the contribution is at most

$$\ll_\alpha \frac{n}{\log n} \sum_{n^{1-2\alpha} < p < q \le n^\alpha} \frac{1}{pq} = O_\alpha\left(\frac{n}{\log n}\right),$$

because the reciprocal prime sum over a fixed-power interval is $O_\alpha(1)$. Thus $T_\alpha(n) = O_\alpha(n/\log n)$.

After the prime phase, every prime $p \le y$ with $p \notin B$ has either been played as the singleton p or has had some proper multiple played; in either case, no future legal Prolonger move can contain p as a factor. Therefore any remaining composite Prolonger move is supported on B.

We prove that the priority list resolves all primes at most y. After the prime and square phases, any remaining legal move supported on B must be squarefree. If it has support size 1, it is comparable with a played or blocked prime. If it has support size 2, then the corresponding semiprime was either played in phase three or was already illegal by the skipped-move invariant. If it has support size 3, the same statement holds for phase four, where the invariant rules out proper multiples.

If a remaining legal move x had support size at least 4 inside B, choose three primes p, q, r from its support. The triple pqr was addressed in phase four. If it was played, then $pqr \mid x$ and x is illegal. If it was not legal when considered, then some earlier played move comparable with pqr must divide pqr or equal pqr by the skipped-move invariant; since all earlier non-prime moves in the priority list have support size at most 3 and are squarefree except the squares p^2, that played divisor also divides x. Again x is illegal. Thus no remaining legal move contains a prime at most y.

As in the proof of Proposition A.10, once all primes at most y are resolved, the residual legal set has size $O_\alpha(n/\log n)$. Combining this residual bound with the phase counts proves the proposition. □

Appendix B. Numerical Computations

This appendix records the numerical computations accompanying the main Bonferroni–4 constant. The computations certify the quoted constants and stress-test the finite combinatorial steps, but they are not substitutes for the analytic proof in Section 7.

B.1. **The Bonferroni–4 constant.** The main theorem uses the following certified upper bound.

Proposition B.1 (Interval-arithmetic certification of $\mathcal{W}_4/2 < 0.19$)**.** *For the moment constants in Section 7,*

$$\begin{aligned} J_1 &\in [0.7885305658,\ 0.7885305661],\\ J_2 &\in [0.18681848,\ 0.18682451],\\ J_3 &\in [0.02009209,\ 0.02009370],\\ J_4 &\in [0.00122263,\ 0.00122282]. \end{aligned}$$

Consequently

$$\frac{\mathcal{W}_4}{2} \le 0.1897123371 < 0.19.$$

In particular one may take a numerical margin $\eta = 2.87 \cdot 10^{-4}$ in $\mathcal{W}_4/2 \le 0.19 - \eta$.

Proof. Let $\mu(du) = \rho(u)\,du$. Then

$$r!J_r = \mu^{\otimes r}\{(u_1,\dots,u_r) : u_1 + \cdots + u_r \le 1\}.$$

Take $N = 10^5$ and partition $(0,1]$ into

$$I_i = \left(\frac{i-1}{N}, \frac{i}{N}\right] \qquad (1 \le i \le N).$$

Write $m_i = \mu(I_i)$. These masses are computed from an exact antiderivative: if $x \in (1/(h+1), 1/h]$, then

$$F(x) := \int_0^x \rho(u)\,du = T_{h+1} + \frac{1}{h+1}\log((h+1)x),$$

where

$$T_k := \sum_{\ell=k}^{\infty} \frac{\log(1+1/\ell)}{\ell+1}.$$

Thus $m_i = F(i/N) - F((i-1)/N)$. For certification the tails T_k are truncated at $H^\star = 10^6$. Since

$$\frac{1}{(\ell+1)^2} \le \frac{\log(1+1/\ell)}{\ell+1} \le \frac{1}{\ell(\ell+1)},$$

we have

$$\sum_{\ell > H^\star} \frac{\log(1+1/\ell)}{\ell+1} \le \sum_{\ell > H^\star} \frac{1}{\ell(\ell+1)} = \frac{1}{H^\star+1}.$$

All logarithms, sums, differences, and products are evaluated with outward-rounded interval arithmetic. A reference Python implementation using `mpmath.iv` for the interval cell masses is available at `scripts/wfour_certification.py` in the artifact repository [Bud26]. Running it with the default parameters computes the cell masses by splitting grid cells at the exact breakpoints $1/h$, performs the four interval convolutions with an explicit FFT round-off enclosure, and reproduces intervals contained in the certified intervals displayed above.

Let $c_s^{(r)}$ be the interval convolution coefficients

$$c_s^{(r)} := \sum_{i_1+\cdots+i_r=s} m_{i_1}\cdots m_{i_r}.$$

Every product cell with $i_1+\cdots+i_r \le N$ lies wholly inside the simplex, while every product cell which can meet the simplex satisfies $(i_1-1)+\cdots+(i_r-1)\le N$, equivalently $i_1+\cdots+i_r \le N+r$. Therefore

$$\frac{1}{r!}\sum_{s\le N} c_s^{(r)} \le J_r \le \frac{1}{r!}\sum_{s\le N+r} c_s^{(r)}.$$

The upper condition $s \le N+r$ is a harmless overcount of one convolution shell relative to the exact meeting condition $s \le N+r-1$; the displayed upper bounds are therefore slightly looser than sharp but remain rigorous. Carrying out these outward-rounded interval convolutions gives the four intervals displayed above. Finally, using the lower endpoints for J_1, J_3 and the upper endpoints for J_2, J_4,

$$\frac{\mathcal{W}_4}{2} \le \frac{1}{2}(1-0.7885305658+0.18682451-0.02009209+0.00122282) = 0.1897123371 < 0.19.$$

□

B.2. **Directed-rounding implementation details.** The certification in Proposition B.1 uses only standard outward-rounded interval operations. The implementation may use IEEE-754 directed rounding modes, or a library with documented correctly rounded elementary functions. For each positive input x, compute lower and upper logarithms

$$\log_{\mathrm{dn}}(x) \le \log x \le \log_{\mathrm{up}}(x).$$

All lower partial sums are accumulated with round-down arithmetic, and all upper partial sums with round-up arithmetic. If an interval antiderivative value is written $F(x) \in [F_{\mathrm{dn}}(x), F_{\mathrm{up}}(x)]$, then the cell mass $m_i = F(i/N) - F((i-1)/N)$ is enclosed by

$$m_i^- = \mathrm{round}_{\mathrm{dn}}\big(F_{\mathrm{dn}}(i/N) - F_{\mathrm{up}}((i-1)/N)\big),$$
$$m_i^+ = \mathrm{round}_{\mathrm{up}}\big(F_{\mathrm{up}}(i/N) - F_{\mathrm{dn}}((i-1)/N)\big).$$

The convolution coefficients $c_s^{(r)}$ are then enclosed by interval polynomial multiplication. A direct implementation may round each product and each accumulated partial sum outward. The reference script uses an FFT accelerator, but encloses its floating-point round-off by an explicit ℓ^1-error bound for each convolution and propagates that bound through the later products. Concretely, for a Cooley–Tukey FFT of transform length L and nonnegative inputs a, b with ℓ^1-norms $\|a\|_1, \|b\|_1$, the script bounds the total ℓ^1 round-off error in the computed $a * b$ by

$$256 \cdot \varepsilon_{\mathrm{mach}} \cdot (\lceil \log_2 L \rceil + 1)^2 \cdot \|a\|_1 \|b\|_1, \tag{B.1}$$

where $\varepsilon_{\mathrm{mach}}$ is the IEEE-754 double-precision machine epsilon. This form dominates the standard γ_k error model of Higham [Hig02, Chapter 24] applied to the two forward transforms, the pointwise multiplication, and the inverse transform; the constant 256 is an intentionally conservative safety factor. The bound (B.1) is added outward to the computed convolution, and the widened interval is used in all downstream products. Thus the displayed certification does not rely on an unbounded floating-point FFT error model.

A direct grid computation of the same integrals gives the convergence displayed in Table 1. The theorem only needs the certified strict margin below 0.19, while the table records an independent convergence check.

$$J_r = \frac{1}{r!}\int_{\substack{u_1+\cdots+u_r\le 1\\ 0<u_i\le 1}} \prod_{i=1}^{r} \rho(u_i)\, du_1\cdots du_r$$

is the quantity being approximated.

Grid	J_1	J_2	J_3	J_4	$\mathcal{W}_4$	$\mathcal{W}_4/2$
2^{13}	0.78844185	0.18676003	0.02008041	0.00122154	0.37945930	0.18972965
2^{14}	0.78849100	0.18679383	0.02008737	0.00122220	0.37943767	0.18971883
2^{15}	0.78851725	0.18681153	0.02009097	0.00122255	0.37942586	0.18971293
2^{16}	0.78852127	0.18681405	0.02009149	0.00122260	0.37942388	0.18971194
2^{17}	0.78852704	0.18681777	0.02009220	0.00122266	0.37942118	0.18971059

TABLE 1. Grid convergence for the Bonferroni–4 truncation.

B.3. **Finite sanity checks.** Four finite checks were used to guard against implementation or formulation errors in the proof.

(1) *Odd-part injection.* Exhaustive verification at $n = 14$ over all 732 antichains found no collision under the odd-part map.
(2) *Monotone replacement.* Exact inclusion–exclusion counts at $n = 10^4, 10^5, 10^6$ for test prime pairs (q_j) and (p_j) with $q_j \le p_j$ confirmed $N(q_1, \dots, q_K) \le N(p_1, \dots, p_K)$.
(3) *Squarefree divisor counts.* For $r \le 4$, the relevant squarefree r-fold divisor counts match the scale $O_r(n(\log\log n)^{r-1}/\log n)$ used in the Bonferroni floor-error bound.
(4) *Controlled prime rounding.* Synthetic monotone comparison prefixes were rounded upward to increasing odd-prime prefixes at $n = 10^4, 10^5, 10^6, 10^8$; the resulting changes in the moment terms were small and consistent with the prime-rounding bridge.

Check	Recorded outcome
Odd-part injection	`(True, None, None, 732)` at $n = 14$
Monotone replacement	Verified at $n = 10^4, 10^5, 10^6$ for listed test prefixes
D_r scaling	Ratified for $r = 1, 2, 3, 4$ up to $n = 10^6$
Prime-rounding bridge	Tested on synthetic prefixes up to $n = 10^8$

TABLE 2. Finite checks for the combinatorial components of the upper-bound proof.

B.4. **Interpretation.** The numerical artifacts support three claims: the constant $\mathcal{W}_4/2 < 0.19$ is reproducible; the finite steps in Theorems 7.10 and 7.20 behave as expected in representative regimes; and the sharper limiting value near 0.18969 is numerically stable. The theorem in this paper is the finite Bonferroni–4 bound below 0.19, not a claim that the full alternating series has been analytically justified.

B.5. **Reproducibility.** From the repository root, the certificate is reproduced by running

```
python3 scripts/wfour_certification.py.
```

The default parameters are $N = 10^5$, $H^\star = 10^6$, and 80 decimal digits; no flags are needed to reproduce intervals contained in the displayed certified intervals. The script imports only the Python standard library, `mpmath`, and `numpy`; the run used for this manuscript was tested with Python 3.14.3, `mpmath` 1.3.0, and `numpy` 2.4.2 on macOS Darwin 25.2.0 (arm64). No separate stdout transcript is archived in the repository; running the command prints a self-describing block containing the four J_r intervals, the final $\mathcal{W}_4/2$ interval, the margin to 0.19, and the elapsed wall-clock time.

APPENDIX C. FORMALIZATION ARTIFACT MAP

This appendix records the status of the Lean artifacts corresponding to the paper's principal results. The map distinguishes between fully formalized finite or structural cores and prose components which use standard analytic number theory or conditional game-theoretic hypotheses. For all entries, verification was last recorded at artifact repository tag `arxiv-v1` (see [Bud26]) and Lean toolchain `leanprover/lean4:v4.28.0`.

C.1. **Verification status by result.**

Result	Status	Lean file(s)	Notes
Theorem 3.1	Formally verified	`shield_reduction/.../ShieldReduction.lean`	Zero-sorry finite divisibility-poset proof.
Theorem 4.1	Partially formalized	`theorem_A/.../ShieldDefs.lean`; `ShieldBasicLemmas.lean`; `ShieldMainTheorem.lean`	Structural identities Lean; analytic inputs (Mertens, PNT) and assembly are prose.
Theorem 4.4	Prose only	None	Unconditional game lower bound proved in the text; no Lean or numerical artifact is claimed.
Theorem 4.7	Conditional; finite cores Lean-verified	`T2Finite/GraphGame.lean`; `T2Finite/HyperGame.lean`; `T2Finite/ResidualComparison.lean`	Conditional on the restricted safe-edge hypothesis. The safe-edge hypothesis and asymptotic activation wrapper remain prose.
Theorem 5.1	Formally verified	`tau_5_24/.../Cover.lean`; `Packing.lean`; `Tau.lean`	Cover and packing identities are zero-sorry Lean; the final floor-sum asymptotic is elementary prose.
Theorem 6.2	Formally verified	`shortener_13_36/.../Compression.lean`; `Sieve.lean`; `MainTheorem.lean`	Zero-sorry odd-part compression and Bonferroni core; standard prime-prefix estimates are used in the prose setup.
Theorem 6.3	Partially formalized	`shortener_5_16/.../Defs.lean`; `Theorems.lean`	Compression, sieve, and algebraic optimization are Lean-checked; one placeholder bundles Chebyshev prime-prefix input and the game-tree prefix wrapper.
Theorem 7.21	Partially formalized	`Round15Bonferroni4/Target.lean`	Lean covers the endgame reduction in `Round15Bonferroni4/Target.lean`; envelope, inversion, measure convergence, and prime-rounding bridge are prose.
Theorem 7.20	Prose only	None	The queued prime-rounding bridge is proved in Section 7; no separate Lean artifact is claimed.
Proposition B.1	Prose with numerical stress testing	`scripts/wfour_certification.py`	Outward-rounded interval certificate computes the four displayed J_r intervals and the final margin.
Proposition 8.1	Prose with numerical stress testing	`scripts/sa_barrier_verification.py`	Proof is prose; the script exhaustively checks the lifted covering inequalities for small N with exact rational arithmetic.
Proposition 8.2	Prose only	None	Extremal shadow argument is presented in the text; no artifact is claimed.
Proposition 8.3	Prose only	None	Arithmetic closure limitation is presented in the text; no artifact is claimed.

Shield Reduction, Theorem 3.1.

Path: `erdos-872/lean/shield_reduction/shield_reduction_aristotle/RequestProject/ShieldReduction.lean`. Declarations: `shield_reduction`. Status: zero-sorry Lean verification; no remaining assumptions beyond the finite definitions of the divisibility poset.

Exact $5/24$ **cover, Theorem 5.1.**

Path: `erdos-872/lean/tau_5_24/tau_5_24_aristotle/RequestProject/Tau.lean`. Declarations: `coverSet_is_cover` and `cover_card_lower_bound`. Status: the structural cover and packing identities are zero-sorry Lean-verified; the asymptotic floor-sum count $5n/24 + O(1)$ is an elementary prose calculation.

$13/36$ **upper bound, Theorem 6.2.**

Path: `erdos-872/lean/shortener_13_36/shortener_13_36_v2_aristotle/RequestProject/Shortener/MainTheorem.lean`. Declarations: `antichain_dfree_bound` and `main_upper_bound`. Status: the odd-part compression and Bonferroni core are zero-sorry Lean-verified; the paper-level exposition uses standard prime-counting estimates for the prime-prefix setup.

Polynomial shield-weight lower bound, Theorem 4.1.
Path: `erdos-872/lean/theorem_A/theorem_A_shield_lower_bound_aristotle/RequestProject/ShieldMainTheorem.lean`. Declarations: `shield_lower_bound` and `barrier_exponent`. Status: the combinatorial core, including the prime-antichain construction, log-budget argument, and exchange principle, has a Lean artifact. Remaining sorries are Mertens' theorem, a Chebyshev-type prime-counting bound, a Chebyshev-ϑ asymptotic, and the final assembly using those analytic inputs.

5/16 **upper bound, Theorem 6.3.**
Path: `erdos-872/lean/shortener_5_16/shortener_5_16_aristotle/RequestProject/Shortener516/Theorems.lean`. Declarations: `antichain_three_five_core_bound`, `game_value_per_parameter`, and `main_theorem`. Status: the compression, sieve, and algebraic optimization core has a Lean artifact. One remaining bundled placeholder covers the Chebyshev prime-prefix estimate and the game-tree wrapper realizing the prescribed prefix.

Conditional T2 lower bound, Theorem 4.7.
Paths:

`erdos-872/lean/erdos_872_core/RequestProject/T2Finite/GraphGame.lean`
`erdos-872/lean/erdos_872_core/RequestProject/T2Finite/HyperGame.lean`
`erdos-872/lean/erdos_872_core/RequestProject/T2Finite/ResidualComparison.lean`

Declarations:

`Q8_maker_ge_add_selfPot`
`HQ8_maker_ge_add_selfPot_of_count_le_two`
`comparison_package`

Status: the finite graph-potential core, scored hypergraph core, and local arithmetic embedding core are Lean-verified. The restricted safe-edge hypothesis is conditional, and the asymptotic activation wrapper remains in prose.

Main $0.19n$ **upper bound, Theorem 7.21.**
Path:

`erdos-872/lean/erdos_872_core/RequestProject/Round15Bonferroni4/Target.lean`

Declaration:

`eventually_strict_lt_point19_of_componentwise_close`

Status: the endgame reduction from sufficiently close first-four moment values to the strict inequality $L(n) < 0.19n$ is zero-sorry Lean-verified. The local density theorem, monotone envelope and inversion, and prime-rounding bridge are the prose moment-convergence inputs proved in Section 7.

This map should be read conservatively. A zero-sorry finite core is not being advertised as an end-to-end verification of every analytic asymptotic in the same theorem unless the entry says so explicitly. Conversely, when the remaining obligations are standard analytic-number-theory inputs, the novelty of the combinatorial argument is already isolated in the formal artifact.

Artifact Availability

The public artifact repository for this paper is available at `https://github.com/xa8zz/erdos-harness`; the version corresponding to this submission is tagged `arxiv-v1` in the repository, and [Bud26] records the artifact metadata. It contains the Lean formalizations, Aristotle submission files and outputs, numerical scripts, and supporting research notes used in the preparation of the present manuscript. Repository-relative paths quoted in Appendices B and C are relative to that artifact root.

Acknowledgments

The author thanks Liam Price, Abhimanyu Adenwalla, StijnC, natso26, Xiao_Hu, and Desmond Weisenberg for the forum refinement chain leading from the first linear bound to the previous public record, and Thomas Bloom for maintaining `erdosproblems.com`. The author also wishes to record substantial gratitude for the research assistance provided by OpenAI's GPT-5.4 Pro, Anthropic's Claude Opus 4.7, and Google's Gemini 3.1 DeepThink, coordinated through a custom multi-round research harness built by the author. GPT-5.4 Pro produced the majority of the exploratory mathematical output used in this project, with Claude and Gemini providing extensive auditing, synthesis, and drafting support. A separate methodological paper and public-facing writeup will describe that workflow in detail.

Declaration of Generative AI and AI-assisted Technologies in the Research and Writing Process

During the preparation of this work, the mathematical content was produced primarily by generative AI systems—OpenAI GPT-5.4 Pro, Anthropic Claude Opus 4.7, and Google Gemini 3.1 DeepThink—coordinated through a custom multi-round research harness designed and operated by the author. GPT-5.4 Pro generated the majority of the novel mathematical arguments, including the shield reduction, the polynomial shield barrier, the piecewise-density analysis underlying the fourth-order Bonferroni argument, and the proof-class obstruction taxonomy. Claude Opus 4.7 and Gemini 3.1 DeepThink provided adversarial auditing, cross-model verification, synthesis across rounds, and drafting support. The author's contributions were the harness design, problem selection and scoping, round-by-round direction and adversarial prompting, selection and integration of outputs into a coherent manuscript, Lean/Aristotle formalization work, numerical verification, and final writing. The author takes full responsibility for the correctness of the claims, the validity of the proofs as presented, the accuracy of the citations, and the integrity of the formal artifacts. A description of the harness and workflow is available at https://www.sensho.xyz/autonomous-research.

References

[And87] Ian Anderson, *Combinatorics of finite sets*, Oxford University Press, 1987.

[BHW16] Csaba Biró, Paul Horn, and D. Jacob Wildstrom, *An upper bound on the extremal version of hajnal's triangle-free game*, Discrete Applied Mathematics **198** (2016), 20–28.

[Blo26] Thomas F. Bloom, *Erdős problem #872*, https://www.erdosproblems.com/872, 2026, Accessed 2026-04-20.

[Bud26] Om Buddhdev, *erdos-harness: Artifacts and formalization files for erdős problem #872*, https://github.com/xa8zz/erdos-harness, 2026, Public artifact repository for this paper; the version corresponding to this submission is tagged `arxiv-v1` in the repository.

[CE90] Peter J. Cameron and Paul Erdős, *On the number of sets of integers with various properties*, Number Theory (Banff, AB, 1988), de Gruyter, Berlin, 1990, pp. 61–79.

[Erd92] Paul Erdős, *Some of my forgotten problems in number theory*, Hardy-Ramanujan Journal **15** (1992), 34–50.

[Erd26] Erdős Problems Forum contributors, *Discussion thread for erdős problem #872*, https://www.erdosproblems.com/forum/thread/872, 2026, Includes refinements by Adenwalla, StijnC, natso26, Xiao_Hu, Desmond Weisenberg, and others, leading to the bound $(419/1008 + o(1))n$; accessed 2026-04-20.

[FS92] Zoltán Füredi and Ákos Seress, *On hajnal's triangle-free game*, Graphs and Combinatorics **8** (1992), no. 1, 75–79.

[Hig02] Nicholas J. Higham, *Accuracy and stability of numerical algorithms*, 2nd ed., SIAM, Philadelphia, 2002.

[HW08] G. H. Hardy and E. M. Wright, *An introduction to the theory of numbers*, 6 ed., Oxford University Press, 2008.

[Kuc24] Petr Kucheriaviy, *Erdős inequality for primitive sets*, 2024.

[Lic23] Jared Duker Lichtman, *A proof of the erdős primitive set conjecture*, Forum of Mathematics, Pi **11** (2023), e18.

[MP10] Greg Martin and Carl Pomerance, *Primitive sets with large counting functions*, 2010.

[Pri26] Liam Price, *First linear upper bound for erdős problem #872*, https://www.erdosproblems.com/forum/thread/872, 2026, Forum contribution; the author disclosed AI-assisted derivation in the post itself.

[SA90] Hanif D. Sherali and Warren P. Adams, *A hierarchy of relaxations between the continuous and convex hull representations for zero-one programming problems*, SIAM Journal on Discrete Mathematics **3** (1990), no. 3, 411–430.

Independent Researcher
Email address: sensho@sensho.xyz