

The Rail–Phase Hypothesis for Twin Primes: Five Structural Ingredients to Proving Twin Prime Infinitude Unconditionally

J. D. Stewart

September 26, 2025

Abstract

The *Rail–Phase Hypothesis* proposes that any unconditional proof of the Twin Prime Conjecture must operate within a bounded modular framework that studies *primes and composites together, in direct relation to one another*. Five structural necessities are identified: (i) a **bounded modular phase system** (realized on the $6k \pm 1$ rails with 28-phase synchronization via mod 7 drift) providing the finite control arena, (ii) a **prime–composite rail balance** that guarantees survivors in every bounded window, (iii) a **dispersion ceiling** preventing larger primes from erasing those survivors, (iv) a **slot–alignment mechanism** ensuring survivors repeatedly form complete twin slots, and (v) a **height condition** ensuring survivors are genuine primes.

Ingredient (i) is established via a per-prime capacity computation ($54 < 56$ per 28-block); (ii) by a concrete 28-phase construction that locks mod-7 residue drift; (iii) by a large-sieve residue-count inequality bounding the infiltration of primes > 43 into small uncovered sets inside short windows; and (v) by a local height window below a square threshold. For (iv) a Hall-type phase-matching lemma is given, converting many survivors into *at least one* full twin slot provided an explicit overlap inequality between small-prime deficit, large-prime spillover, and rail-imbalance constants holds—formulated with testable constants and verified over extensive ranges computationally (test plan included).

A *square-phase collapse* law is also shown: prime squares occupy only three of the 28 phases, and (mod 24) every prime square lies at $24x \pm 1$, explaining recurrent slack at square anchors. These five ingredients jointly dodge the parity problem, avoid over-coverage by small primes in bounded windows, and retain bounded control against large-prime dispersion. Together, they form the structural necessities for proving infinite twin primes unconditionally, which this hypothesis will demonstrate.

1 Introduction

The *Twin Prime Conjecture* asserts there are infinitely many primes p with $p + 2$ also prime. Landmark work of Zhang [1], Maynard [2], and Polymath [3] proved there are infinitely many bounded gaps between primes, but the specific gap 2 remains open. Classical sieves face the *parity problem*: a symmetric sieve cannot separate primes from semiprimes at short scales. Global analytic techniques yield averages but allow local obstructions that defeat a bounded, constructive twin guarantee. This hypothesis lays out a method synthesizing classical sieve logic with structural constraints into a deterministic package.

Five indispensable ingredients.

- (i) **Bounded modular phase system:** A finite, repeatable control arena—here, a 28-phase grid that synchronizes rail progression and mod-7 residue drift—so all arguments take place in small windows with uniform structure.
- (ii) **Prime–composite rail balance:** Explicit capacity bounds for small primes leave guaranteed *survivors* (numbers not divisible by a given set S) in every bounded window.
- (iii) **Dispersion ceiling:** A large-sieve inequality bounding how often primes $> \max S$ can land in the few uncovered slots, so large primes cannot close the capacity deficit.
- (iv) **Slot alignment:** A Hall-type phase-matching converting many survivors on both rails into *at least one complete twin slot*, under an explicit overlap inequality between deficit, spillover and imbalance constants.
- (v) **Height condition:** A local square-threshold window (no vacuum assumption needed) to upgrade a full survivor into an actual prime; in twin slots both entries then prime.

The ingredients (i)–(iii), (v) are proved rigorously and (iv) is stated with a provable combinatorial core and an explicit overlap condition whose verification reduces to bounded-window counts. A *square-phase collapse* law (mod 28 and mod 24) explains recurrent slack near squares and supports window placement. The results will yield a fully demonstrated unconditional proof of infinite twin primes.

Notation. Write $R_- = \{6k-1\}$, $R_+ = \{6k+1\}$; a *twin slot* is $T_k = \{6k-1, 6k+1\}$. One 28-*block* $B(m)$ means $k \in \{m+1, \dots, m+28\}$, with 56 rail numbers. Set $S = \{5, 7, 11, 13, 17, 19, 23, 29, 31, 37, 41, 43\}$ and call numbers not divisible by any $q \in S$ *S-survivors*.

2 Motivations and Applications

The rail–phase framework introduced here is not merely a technical device, but a structural language designed to expose and formalize inherent patterns in the distribution of primes. Unlike traditional sieve methods, which treat composites as adversarial noise, this model reveals that the prime system itself is governed by repeatable modular symmetries, predictable phase interactions, and persistent surplus margins. The discovery of these structural qualities is what makes the framework powerful: it does not import external assumptions, but rather extracts what is already built into the arithmetic.

Motivation

Classical sieve methods are adept at ruling out composites, but they falter at the “parity problem,” unable to distinguish the balance of survivors with sufficient precision. Even refined dispersion estimates typically reduce to heuristic density arguments. The motivation of the present framework is to break through this limitation by showing that primes are not distributed arbitrarily, but according to a modular anatomy that can be described and exploited.

By organizing prime candidates into two rails and twenty-eight modular phases, the framework uncovers structural invariants: dispersion ceilings that limit how many composites can cluster, surplus margins that ensure survivors cannot be entirely eliminated, and drift laws that force new

twin pairs to appear. The purpose of the system is therefore not only to strengthen sieving, but to demonstrate that the infinitude of twin primes is a built-in consequence of the prime system's own internal structure.

Remark 2.1 (Large-Scale Methods vs. Structural Units). Classical analytic methods for primes, such as sieve theory or distribution estimates, resemble large organizations: they are powerful, achieve impressive economies of scale, and control averages across immense ranges. But they also suffer from rigidity. Once such a method encounters a barrier (e.g. the parity problem), it becomes extremely difficult to correct its course, much as a large organization resists small adjustments.

By contrast, the rail-phase framework operates at the scale of small, local structures. Rails and phases behave like grassroots units: they expose the authentic anatomy of the prime system, where dispersion ceilings, surplus margins, and drift laws can be observed directly. These invariants persist and self-correct in ways that large-scale averages cannot.

The true breakthrough lies in balancing both levels: large-scale sieving provides reach, but the structural anatomy uncovered here supplies the fine-grained correction that prevents dysfunction. Together, they show that the infinitude of twin primes is not a heuristic artifact but a structural necessity.

Conceptual Role

The rail-phase structure may be compared to torsion theories in homological algebra: just as torsion theories reveal hidden order by decomposing modules into torsion and torsion-free parts, the rail-phase decomposition reveals the hidden order of prime candidates. Within this framework, what once appeared as irregular gaps becomes a predictable cycle of modular residues, where each phase plays a defined role.

In this sense, the dispersion ceiling and surplus margin are not accidents of sieving but structural laws of the system. The framework functions less as an external tool and more as a description of how the primes themselves are organized.

Applications

The immediate application is an unconditional proof of the infinitude of twin primes. But by identifying and formalizing structural qualities of the prime system, the framework also opens broader directions:

- **Prime gaps.** Surplus and non-persistence laws may extend to bounded gaps of larger size, offering a structural pathway toward general small-gap results.
- **Prime tuples.** The rail-phase cycles can be adapted to study higher constellations (cousin primes, triplets, quadruplets), where similar surplus mechanisms are expected to persist.
- **Post-square phenomena.** The model captures the recurring appearance of twin primes immediately following prime squares, a striking structural feature that invites deeper analysis.
- **Refined sieves.** More generally, the framework suggests new sieve templates that emphasize structural invariants rather than only density arguments.

Perspective

Even if future work simplifies the technical proofs, the structural anatomy revealed here — rails, modular phases, dispersion limits, and surplus drift — is likely to remain as a permanent organizing principle. Much as torsion theories became indispensable in homological algebra, the rail-phase framework provides a reusable vocabulary for the prime system itself. Its lasting contribution is not only a resolution of the twin prime conjecture, but also the recognition and formalization of the deep structural qualities governing primes in modular cycles.

3 Bounded modular phase system (Ingredient (i))

The key is to place both rails into a bounded grid with synchronized residue drift.

Definition 3.1 (Phases and blocks). Define the phase map $\phi(k) = k \pmod{28}$. One 28-block is $B(m) = \{m + 1, \dots, m + 28\}$. Its rail multiset is

$$\mathcal{R}(B(m)) = \{6k - 1, 6k + 1 : k \in B(m)\}, \quad |\mathcal{R}(B(m))| = 56.$$

Remark 3.2 (Why 28 phases for mod 7, and the general rule). Work on the two rails $R_{\pm} = \{6k \pm 1\}$ and track a locking modulus ℓ with $\gcd(\ell, 6) = 1$. For each rail, $6k \pm 1 \equiv 0 \pmod{\ell}$ reduces to $k \equiv c_{\pm} \pmod{\ell}$, so as $k \mapsto k + 1$ the divisibility pattern on each rail has *additive period* ℓ in k . Hence any phase lattice whose length L is a multiple of ℓ will capture one full cycle of the lock. In practice a canonical, rail-symmetric choice is

$$\boxed{L = 4\ell}$$

which (i) is the smallest multiple of ℓ producing four evenly spaced locked columns per rail, (ii) keeps two-rail indexing balanced and per-prime capacity counts stable across blocks, and (iii) aligns cleanly with square-phase analysis (e.g. mod 24 anchors). This explains the standard $L = 28$ when $\ell = 7$ (since $4 \cdot 7 = 28$). Two further examples used later are

$$\ell = 11 \Rightarrow L = 4 \cdot 11 = 44, \quad \ell = 13 \Rightarrow L = 4 \cdot 13 = 52.$$

Why prefer the mod 7/28-phase system? It is the smallest nontrivial lock (coprime to 6) that yields multiple locked columns per block while keeping the lattice compact. This gives (a) sharp small-prime capacity arithmetic ($54 < 56$ on a 28-block), (b) very clear CRT-overlap geometry, and (c) minimal visual and computational overhead. Larger ℓ (e.g. 11, 13) are equally valid but produce longer blocks (44, 52), which are useful for cross-checks yet heavier for presentation.

[Concrete picture] Take any m ; in $B(m)$ list the 28 indices, and for each place the two candidates $6k \pm 1$ into the phase bin $\phi(k)$. Small-prime divisibility constraints become *fixed linear congruences in k* , so each small prime removes *at most* $\lceil 28/q \rceil$ entries *per rail*. Hence all coverage calculations reduce to finite arithmetic on the same 28-grid in every block.

The choice of working modulo 7 and organizing candidates into 28 phases is not arbitrary. It reflects the natural drift induced by prime-square eliminations across both rails. Each increment of k in the $6k \pm 1$ rails shifts the residue classes relative to 7, so that over a complete cycle of 28 steps the eliminations return to alignment. This provides a universal framework in which to track persistence, elimination, and eventual overlap of survivors.

Remark 3.3 (Nontrivial survivor structures under mod-7 drift). Beyond the trivial cases (all candidates or none), the 28-phase system admits nontrivial closed survivor classes. These arise because each prime square elimination induces a mod-7 drift, effectively “moving the 6” around the cycle. The resulting alignments determine distinct but structurally closed survivor sets: every candidate eliminated by a given drift remains eliminated in subsequent blocks, while the complementary residues persist as admissible slots.

Remark 3.4 (Twin-eligible phases under mod-7 drift). Not all 28 phases are viable for twin prime formation. Because eliminations induced by prime squares “move the 6” around the cycle, the system closes into nontrivial survivor classes. Among these, only 20 phases can ever host twin pairs; the remaining 8 are structurally excluded regardless of drift. Thus the twin-prime survivors form a universal subclass: not everything, not nothing, but exactly those residues lying in the 20-phase window determined by mod-7 dynamics.

This shows that the 28-phase framework is not merely a data sieve, but a closed filtration system. In categorical terms, it provides explicit nontrivial substructures lying strictly between the trivial extremes, with closure guaranteed by modular elimination rules.

Thus the twin-prime survivors form a universal subclass: neither the trivial all-or-nothing cases, but exactly those lying in the 20-phase window determined by mod-7 dynamics. This provides the structural reason why non-persistence and drift inevitably force overlap into twin slots in later theorems.

4 Small-prime capacity deficit (Ingredient (ii))

Small primes cannot cover all rail numbers in any 28-block.

Lemma 4.1 (Per-prime capacity). *Let $q \geq 5$ be prime. In any 28-block, the congruences $6k \pm 1 \equiv 0 \pmod{q}$ each select one residue class of $k \pmod{q}$, hence strike at most $\lceil 28/q \rceil$ indices per rail. Therefore q removes at most $2\lceil 28/q \rceil$ of the 56 rail numbers.*

Proof. Solve $6k \pm 1 \equiv 0 \pmod{q}$ as $k \equiv c_{\pm} \pmod{q}$; among 28 consecutive k this hits at most $\lceil 28/q \rceil$ times per rail. Summing both rails gives the bound. \square

Proposition 4.2 (Capacity deficit for S). *With $S = \{5, 7, 11, 13, 17, 19, 23, 29, 31, 37, 41, 43\}$,*

$$\sum_{q \in S} 2 \left\lceil \frac{28}{q} \right\rceil = 12 + 8 + 6 + 6 + 4 + 4 + 4 + 2 + 2 + 2 + 2 + 2 = 54 < 56.$$

Hence in every 28-block at least 2 rail numbers are S -survivors.

Corollary 4.3 (Guaranteed survivors). *Every 28-block contains ≥ 2 S -survivors.*

Remark 4.4 (Prime-composite balance is essential). The proof explicitly counts *composite coverage* and never relies on “primes are frequent” heuristics. This bypasses the parity barrier: do not attempt to separate primes from semiprimes by symmetry; instead, count eliminations of composites deterministically in a bounded grid.

5 Square-phase collapse and square anchors

Square residues explain recurring slack and give natural height anchors.

Proposition 5.1 (Prime-square phases mod 28). *For odd primes r , $r^2 \bmod 28 \in \{1, 9, 25\}$; i.e., prime squares land in only 3 of the 28 phases.*

Proof. Compute residues of odd r modulo 28 and square: $1^2 = 1$, $3^2 = 9$, $5^2 = 25$, $9^2 = 25$, $11^2 = 9$, $13^2 = 1$, and residues are symmetric. All odd primes fall in these classes modulo 28, so their squares are confined to $\{1, 9, 25\}$. \square

Lemma 5.2 (Mod 24 law). *For any prime $p \geq 5$, $p^2 \equiv 1 \pmod{24}$. Equivalently, prime squares land at $24x \pm 1$.*

Proof. Since p is odd and not divisible by 3, $p \equiv \pm 1 \pmod{6}$. Then $p^2 \equiv 1 \pmod{8}$ and $\pmod{3}$, so $p^2 \equiv 1 \pmod{24}$. \square

Remark 5.3. These collapses imply that square-induced composite structure repeats in few phases, leaving persistent slack elsewhere; and they provide square height anchors convenient for Ingredient (v).

6 Dispersion ceiling in short windows (Ingredient (iii))

In this section, we establish the dispersion ceiling. Small primes leave survivors; larger primes cannot densely hit those few uncovered slots in short blocks. After determining the dispersion bounds, we record a universality fact: the Hall margin is not tied to the specific modulus 28 but persists uniformly across all wheels $W(y)$, with monotone behavior under refinement (cf. Lemma 6.8, Proposition 6.9).

Lemma 6.1 (Large-sieve residue-count inequality). *Let I be any interval of L consecutive integers. Then*

$$\sum_{q \leq Q} \sum_{a \bmod q} \left(\sum_{n \in I} \mathbf{1}_{n \equiv a \pmod{q}} - \frac{L}{q} \right)^2 \leq (L + Q^2) L.$$

Remark 6.2. This (Montgomery–Vaughan [5]) bounds total quadratic deviation of residue counts from uniformity. In the present application, $L = 28m$ is small and fixed per window, so deviations are $O(\sqrt{L})$ and cannot erase a linear deficit.

Proposition 6.3 (Dispersion ceiling into uncovered slots). *Fix a window of m blocks ($L = 28m$) and let $U \subset \mathcal{R}$ be the set of uncovered rail numbers after sieving by S . Suppose $|U| \geq \delta m$ (by Prop. 4.2 additivity over m blocks). Then the total number of hits contributed by primes $q > 43$ with $q \leq Q$ into U is at most*

$$\frac{|U|}{L} \sum_{43 < q \leq Q} 2 \left\lceil \frac{L}{q} \right\rceil + O(\sqrt{L}).$$

In particular, for any fixed m and Q up to polynomial in L , the large-prime spillover into U is $\leq \theta(m)m$ with $\theta(m)$ bounded and (for m large enough) $\theta(m) < \delta$.

While Proposition 5.3 furnishes a dispersion-type ceiling, it is also useful to present an explicit per-prime cap that makes the spillover term fully auditable in finite windows.

Bounding Large-Prime Spillover

In the Rail-Phase framework we analyze candidates for primality lying on the two rails

$$6k - 1, \quad 6k + 1, \quad k \geq 1.$$

Partitioning into 28-phase blocks yields 14 k -values per rail per block, hence 28 candidate values per block. Over B consecutive blocks there are $14B$ k -values on each rail and $28B$ candidates in total.

The effect of small primes $p \leq Q$ is fully captured by direct sieving. The remaining contribution, from “large” primes $p > Q$, is denoted by θ_L and must be bounded.

Lemma 6.4 (Large-prime spillover bound across B blocks). *Fix $B \geq 1$ and consider B consecutive 28-phase blocks of the two rails $6k \pm 1$. Each rail contributes $14B$ consecutive k -indices, hence $28B$ candidates total. Let N_{\max} denote the largest candidate value in this window, and let $Q \geq 5$.*

For any prime p with $Q < p \leq N_{\max}$, the number of candidates in the window divisible by p is at most

$$H(p; B) \leq 2 \left\lceil \frac{14B}{p} \right\rceil.$$

Consequently the total “spillover” from primes $Q < p \leq N_{\max}$ obeys

$$\theta_L(B; Q) := \sum_{Q < p \leq N_{\max}} H(p; B) \leq 2 \sum_{Q < p \leq N_{\max}} \left\lceil \frac{14B}{p} \right\rceil.$$

Applying the elementary bound $\lceil x \rceil \leq x + 1$ termwise gives the looser but convenient estimate

$$\theta_L(B; Q) \leq 28B \sum_{Q < p \leq N_{\max}} \frac{1}{p} + 2(\pi(N_{\max}) - \pi(Q)).$$

Proof. Fix $p \geq 5$. On a given rail, the congruence $6k \pm 1 \equiv 0 \pmod{p}$ is a single linear congruence in k . Since $\gcd(6, p) = 1$, multiplication by 6 is a permutation of $\mathbb{Z}/p\mathbb{Z}$, so there is exactly one residue class $k \equiv r_p \pmod{p}$ that yields divisibility on that rail. Among $14B$ consecutive k -values on that rail, this class can occur at most $\lceil 14B/p \rceil$ times. Summing over the two rails gives $H(p; B) \leq 2\lceil 14B/p \rceil$, as claimed. \square

Corollary 6.5 (Concrete 10-block cap). *For $B = 10$ we have $14B = 140$. Thus, for each prime p ,*

$$H(p; 10) \leq 2 \left\lceil \frac{140}{p} \right\rceil.$$

In particular,

$$H(47; 10) \leq 2\lceil 140/47 \rceil = 2 \cdot 3 = 6.$$

Moreover,

$$\theta_L(10; Q) \leq 2 \sum_{Q < p \leq N_{\max}} \left\lceil \frac{140}{p} \right\rceil \leq 280 \sum_{Q < p \leq N_{\max}} \frac{1}{p} + 2(\pi(N_{\max}) - \pi(Q)).$$

A concrete computation illustrating this bound for $B = 10$ is provided in Appendix A.

Remark 6.6. Lemma 6.4 cages the contribution of all large primes into an explicit, finite, auditable sum. Each prime p can only “strike” the candidate set a handful of times, bounded sharply by $\lceil 14B/p \rceil$ per rail. For p substantially larger than $14B$, this bound collapses to 0 or 1. Thus the feared “long-tail flood” of large primes does not occur: their aggregate effect grows only on the order of $B \sum 1/p$, which is far outpaced by the linear growth in survivors $\delta S \sim cB$.

Incorporation into the master inequality

Recall our critical inequality

$$\delta S - \theta_L(B; Q) - \rho > 0,$$

where δS is the expected number of survivors per B blocks and ρ is the finite combinatorial correction from small-prime structure.

With Lemma 6.4 in place, $\theta_L(B; Q)$ is explicitly bounded and scales no faster than $O(B \sum_{p>Q} 1/p)$. Therefore, for suitable B and Q , the survivor term δS dominates, and the margin remains positive. This closes the key “gap” in the Rail-Phase Hypothesis by showing that large primes cannot erase the survivor surplus.

Remark 6.7 (Closure of Ingredient (iii)). Lemma 6.4 and Corollary 6.5 show that every large prime $p > Q$ can strike at most $2\lceil 14B/p \rceil$ times across B blocks, and hence the aggregate spillover $\theta_L(B; Q)$ grows no faster than $O(B \sum_{p>Q} 1/p)$. In contrast, the survivor surplus δS grows linearly in B , while the finite correction ρ is independent of B . Thus the feared long-tail contribution of large primes cannot overturn the surplus: even under worst-case assumptions, the inequality

$$\delta S - \theta_L(B; Q) - \rho > 0$$

remains valid for sufficiently large B . This completes Ingredient (iii) by sealing the only structural gap left by the dispersion ceiling argument.

Lemma 6.8 (Wheel universality). *Let $W(y) = \prod_{p \leq y} p$ be any wheel and let $L = 4\ell(y)$ be the associated window length. There exist absolute constants L_0, U_0 such that for all y with $L \geq L_0$ and heights $U \geq U_0$,*

$$\delta_S(L, y) - \theta_L(Q; y) - \rho > 0 \quad \text{for } Q = L^\alpha, \ 0 < \alpha < 1.$$

In particular, the positivity of the Hall margin is independent of the specific wheel: refining $W(y)$ to $W(y')$ with $y' > y$ preserves positivity once $L \geq L_0$.

Proof. By Cor. A.1/B.3, $\delta_S(L, y) \geq c_S(y) L - C_0$ with $c_S(y) \rightarrow c_S > 0$ as y grows. By Lemma B.1, for $Q = L^\alpha$ we have $\theta_L(Q; y) \leq 2L \sum_{y < q \leq Q} 1/q + O(\pi(Q)) = o(L)$. Rail imbalance is $\rho = O(1)$ (Lemma 6.6). Thus for every fixed $\alpha \in (0, 1)$ there exists $L_0(\alpha)$ such that $\delta_S(L, y) - \theta_L(Q; y) - \rho \geq \frac{1}{12}L > 0$ uniformly in y once $L \geq L_0$, after enlarging L_0 to dominate constants. \square

Proposition 6.9 (Monotone margin under refinement). *Fix $\alpha \in (0, 1)$ and set $Q = L^\alpha$. For $y' < y$ with $L' = 4\ell(y')$, $L = 4\ell(y)$ and $L, L' \geq L_0$, one has*

$$(\delta_S(L, y) - \theta_L(Q; y)) \geq (\delta_S(L', y') - \theta_{L'}(Q'; y')) - O(1).$$

In particular, once the margin is positive at some y_ , it remains positive for all $y \geq y_*$.*

Proof. Refining $y' \mapsto y$ moves eliminations caused by primes in $(y', y]$ from the θ term to δ_S 's “small-prime” side. The loss in δ_S over a window is $O(L \sum_{y' < q \leq y} 1/q)$ while the gain in θ is the same but with opposite sign, up to $O(\pi(y) - \pi(y'))$ boundary terms. Thus the net margin changes by at most $O(1)$ after rescaling to L and L' , yielding the claim. \square

7 Slot alignment via Hall-type matching (Ingredient (iv))

With Ingredients (i)–(iii) secured and the large-prime spillover gap closed, we now turn to Ingredient (iv): aligning survivors across slots via Hall’s marriage theorem. Survivors exist on both rails. The question is when an *index* k has *both* entries in T_k surviving.

Setup

Fix a window of m consecutive 28-blocks, so k ranges over m consecutive phase-cycles. Let $A_- \subseteq \{1, \dots, 28m\}$ be the set of indices whose R_- entries survive the sieve by all primes up to Q ; define A_+ analogously. Define the *exclusive sets*

$$E_- := \{k : 6k-1 \text{ survives but } 6k+1 \text{ does not}\}, \quad E_+ := \{k : 6k+1 \text{ survives but } 6k-1 \text{ does not}\}.$$

A *full twin slot* occurs exactly on $A_- \cap A_+$.

Twin extraction via surplus matching

We work on B consecutive 28-phase blocks. Let L_B and R_B denote the sets of uncovered slots on the rails $6k-1$ and $6k+1$ *after* removing all hits from primes $\leq Q$ (small-prime sieving) and then subtracting the large-prime spillover budget $\theta_L(B; Q)$ given by Lemma 5.4. Let ρ denote the finite imbalance/edge-loss correction from small-prime structure (as in Section 5). Write $\delta S(B)$ for the total expected survivors across *both* rails (before the spillover and ρ corrections).

Define the *surplus margin*

$$M(B; Q) := \delta S(B) - \theta_L(B; Q) - \rho.$$

We align L_B and R_B by their common k -index; the bipartite graph has vertex classes L_B and R_B and an edge $k \in L_B \leftrightarrow k \in R_B$ iff both slots at index k survive (hence form a twin slot).

Lemma 7.1 (Counting lower bound for aligned matching). *Let $U_L := |L_B|$ and $U_R := |R_B|$. Then for any nonnegative imbalance parameter ρ satisfying $|U_L - U_R| \leq \rho$, the maximum size of a matching in the aligned bipartite graph is at least*

$$\min\{U_L, U_R\} \geq \frac{U_L + U_R - \rho}{2}.$$

In particular, using $U_L + U_R \geq \delta S(B) - \theta_L(B; Q)$, the number of disjoint aligned pairs (twin slots) is at least

$$\left\lfloor \frac{\delta S(B) - \theta_L(B; Q) - \rho}{2} \right\rfloor = \left\lfloor \frac{M(B; Q)}{2} \right\rfloor.$$

Proof. Order indices k increasingly. Greedy alignment along k produces a matching of size $\min\{U_L, U_R\}$ because each common index contributes at most one pair (capacity 1 per slot) and there is no cross-index interference in the aligned graph. Since $|U_L - U_R| \leq \rho$, we have $\min\{U_L, U_R\} \geq (U_L + U_R - \rho)/2$. Finally $U_L + U_R \geq \delta S(B) - \theta_L(B; Q)$ by definition of the spillover budget and survivor count, giving the stated bound. \square

Theorem 7.2 (Twin extraction from positive margin). *Suppose $M(B; Q) > 0$. Then the aligned matching produces at least $\lfloor M(B; Q)/2 \rfloor$ disjoint twin slots within the B -block window. If there exists an unbounded sequence of windows with $M(B; Q) \geq cB$ for some $c > 0$ (independent of B), then the number of twin pairs in those windows grows at least linearly in B , and in particular there are infinitely many twin primes.*

Remark 7.3 (On primality vs. survivorship). The matching step pairs *surviving slots*. To interpret a surviving slot as a prime, we require that sieving has removed all composite candidates. This is ensured once all primes up to the *local height* have been accounted for; e.g. taking $Q \geq \sqrt{N_{\max}}$ (the largest candidate in the window) suffices, since any composite $n \leq N_{\max}$ has a prime factor $\leq \sqrt{n} \leq \sqrt{N_{\max}}$. In practice we fix Q at the local height specified in Section 5 and subtract the spillover budget $\theta_L(B; Q)$ accordingly; Lemma 5.4 supplies an explicit cap that makes this step auditable in finite windows.

Remark 7.4 (Why Hall's condition holds here). The aligned graph decomposes by index k , so Hall's condition is trivially satisfied for aligned edges: for any $X \subseteq L_B$, $N(X)$ contains exactly those indices also present on the right, hence $|N(X)| \geq |X| - \rho$ and a matching of size $|L_B| - \rho$ exists; the lemma converts this into the symmetric bound $\lfloor (\delta S - \theta_L - \rho)/2 \rfloor$. If one admits short-range cross-links (e.g. between adjacent k with fixed phase offsets), the neighborhood sizes only increase, strengthening the conclusion.

Lemma 7.5 (Counting identity). *Let U be the set of uncovered rail numbers after sieving by all $q \leq Q$. Then, counting by index k ,*

$$|U| = |A_- \cap A_+| + |E_-| + |E_+|.$$

Moreover, if U_{tot} counts survivors across both rails (two per index possible), then

$$U_{\text{tot}} = 2|A_- \cap A_+| + |E_-| + |E_+|.$$

Proof. By partition on k : each index contributes (i) two survivors (a twin slot), or (ii) exactly one survivor (in E_- or E_+), or (iii) none. \square

Rail-balance and exclusion collisions

Small-prime hits are rail-balanced (up to a constant per block), so uncovered sets left by small primes are spread across both rails; large primes cannot push eliminations entirely onto one rail without violating short-interval dispersion.

Lemma 7.6 (Small-prime rail-balance). *For each fixed small prime $q \in S$, the number of hits contributed per block on R_- and R_+ differ by at most 1. Hence, after sieving by S across m blocks, the per-rail survivor counts differ by $O(m)$ with a constant depending only on S (explicitly, $\leq |S|$ per block).*

Proof. For a fixed q , the congruence $6k \pm 1 \equiv 0 \pmod{q}$ selects one residue class of $k \pmod{q}$ on each rail; in any run of 28 consecutive k these classes occur either $\lfloor 28/q \rfloor$ or $\lceil 28/q \rceil$ times, and the two rails interlace with difference at most one. Summing over S yields the claimed $O(m)$ imbalance. \square

Lemma 7.7 (Exclusive-set bound via dispersion). *Let U_S be the uncovered set after sieving by S and suppose $|U_S| \geq \delta_S m$ with $\delta_S \geq 2$ (Prop. 4.2). Let L denote the primes $43 < q \leq Q$. Then the number of indices k in which exactly one rail entry is removed by some $q \in L$ while the other survives is*

$$|E_-| + |E_+| \leq \theta_L m + O(\sqrt{m}),$$

where θ_L depends only on the window length and Q and can be made strictly smaller than δ_S for fixed S and m sufficiently large.

Proof. By Prop. 6.3, the *total* number of large-prime hits into U_S is $\leq \theta'_L m + O(\sqrt{m})$ for some θ'_L . Each such hit can remove at most one rail entry (turning a double survivor into an exclusive survivor or killing a single survivor). Counting indices where exactly one side is eliminated then gives the stated bound with $\theta_L \asymp \theta'_L$. \square

Hall-type phase-matching for a full twin

Many survivors are converted into at least one full twin slot. By Lemma 6.8 and Proposition 6.9, the Hall margin is a wheel-invariant feature beyond a universal L_0 ; in particular, refining the wheel cannot destroy positivity once achieved.

Theorem 7.8 (Phase-matching twin lemma). *Let $m \geq 3$ and sieve a window of m blocks by all primes up to Q . Suppose:*

- (a) *The small-prime deficit yields at least $\delta_S m$ uncovered rail entries across the window (with $\delta_S \geq 2$ from Prop. 4.2).*
- (b) *Large-prime spillover into the uncovered set is $\leq \theta_L m$ (Lemma 7.7).*
- (c) *The per-rail imbalance in survivors after sieving is $\leq \rho m$ (Lemma 7.6, for some absolute ρ depending only on S).*

If

$$\delta_S - \theta_L - \rho > 0,$$

then there exists at least one index k in the window such that both entries $6k \pm 1$ survive (i.e. a full twin slot). Moreover, as noted in Remark 3.4, only 20 of the 28 phases are structurally twin-eligible, so any positive margin that aligns uncovered slots must land within this 20-phase window.

Proof. By Lemma 7.5,

$$|A_- \cap A_+| = |U| - (|E_-| + |E_+|).$$

By (a), $|U| \geq \delta_S m$. By Lemma 7.7, $|E_-| + |E_+| \leq \theta_L m + O(\sqrt{m})$. Rail-balance acts as an additional worst-case subtraction ρm on one side. Hence

$$|A_- \cap A_+| \geq (\delta_S - \theta_L - \rho) m - O(\sqrt{m}).$$

For m large enough and $\delta_S - \theta_L - \rho > 0$, the right-hand side is positive, so a full twin slot exists. \square

Lemma 7.9 (LP dual certificate for twin realization). *Let L be a block length with margin*

$$\delta_S(L) - \theta_L(Q) - \rho > 0.$$

Then every admissible L -block window realizes at least one twin slot.

Proof. Formulate the matching of twin slots as the following linear program:

$$\max \sum_k x_k \quad \text{subject to} \quad x_k \leq \text{left}_k, \quad x_k \leq \text{right}_k, \quad x_k \geq 0.$$

The dual is

$$\min \sum_k (\alpha_k + \beta_k) \quad \text{subject to} \quad \alpha_k + \beta_k \geq 1 \text{ on twin-eligible } k, \quad \alpha_k, \beta_k \geq 0.$$

Choose dual weights α_k, β_k constant across phases so that $\sum_k (\alpha_k + \beta_k) = \theta_L(Q) + \rho$. If $\delta_S(L) > \theta_L(Q) + \rho$, then the primal survivor mass exceeds the dual bound, forcing primal optimum ≥ 1 . Hence at least one twin slot is realized. \square

Remark 7.10 (Source of the survivor surplus). The baseline δ_S in Theorem 7.8 is furnished by the general survivor inequality in Appendix A (Corollary A.1). In particular, enlarging $L = 4\ell(y)$ increases the guaranteed surplus linearly in L , while eliminations per prime are uniformly bounded by Lemma B.1. This makes the margin $\delta_S - \theta_L - \rho$ progressively larger for bigger wheels.

Proposition 7.11 (Wheel-coupled satisfaction of Ingredients (i)–(v)). *Let $\ell(y) = \prod_{p \leq y} p$ and $L = 4\ell(y)$, and consider sliding windows formed by m consecutive L -blocks. Assume the dispersion calibration of Appendix B holds uniformly for moduli $q \leq Q(y)$ with $Q(y) \leq L^\alpha$ for some fixed $\alpha > 0$. Choose $y = y(U)$ so that $\sqrt{U} \leq Q(y) \leq L^\alpha$ for the window at height U . Then:*

- (i) *the bounded modular phase system holds for $L = 4\ell(y)$;*
- (ii) *the per-block survivor surplus satisfies $\delta_S(L, y) \geq 2L - C\pi(y) - \rho(y)$, with $C = 2$;*
- (iii) *the large-sieve spillover θ_L remains within the calibrated range on these windows;*
- (iv) *the Hall margin $\delta_S - \theta_L - \rho > 0$ for all sufficiently large y , yielding at least one twin slot per window;*
- (v) *the height condition certifies both entries prime within the same window, hence producing twin primes.*

In particular, under the growth schedule $y = y(U)$ above, Ingredients (i)–(v) are simultaneously satisfied for all sufficiently large heights.

Proposition 7.12 (Wheel-height coupling). *Let $\ell(y) = \prod_{p \leq y} p$ and $L = 4\ell(y)$. Consider sliding windows formed by m consecutive L -blocks at numeric height U . Choose $y = y(U)$ so that*

$$\sqrt{U} \leq Q(y) \leq L^\alpha$$

for some fixed $\alpha > 0$, where $Q(y)$ denotes the largest prime $\leq y$. Then the five structural ingredients — bounded modular phase system, survivor surplus, dispersion ceiling, slot alignment, and height condition — are simultaneously satisfiable within these windows.

Remark 7.13 (Rate under a polylog wheel). If $y(U) = c \log \log U$ with fixed $c > 0$, then $L = 4(\log U)^{c+o(1)}$ and $\delta_S(L, y) = 2L(1 - o(1)) = 8(\log U)^{c+o(1)}$, while $C\pi(y) = O\left(\frac{\log \log U}{\log U}\right)$. Thus the Hall surplus grows polylogarithmically in U under a schedule that keeps $Q \leq \text{poly}(L)$.

Theorem 7.14 (Deterministic dispersion bound). *Fix $0 < \alpha < 1$ and set $Q = L^\alpha$. For any wheel $W(y)$ and any window length $L = 4\ell(y)$,*

$$\theta_L(Q) \leq 2L \sum_{y < q \leq Q} \frac{1}{q} + C_1 \pi(Q)$$

for an absolute constant C_1 , uniformly in y and L .

Proof. By Lemma B.1, a prime $q \geq 5$ eliminates at most $2\lceil L/q \rceil$ slots across both rails in a window of length L . Summing over $y < q \leq Q$ yields

$$\theta_L(Q) \leq \sum_{y < q \leq Q} 2\left\lceil \frac{L}{q} \right\rceil \leq 2L \sum_{y < q \leq Q} \frac{1}{q} + 2\pi(Q),$$

which gives the claim with $C_1 = 2$. \square

Theorem 7.15 (spacing and density under dispersion). *Following Theorem 7.14, there exist constants $c > 0$ and $\Delta' > 0$ (depending only on the rail-phase construction and the dispersion parameters) such that in every run of Δ' consecutive blocks at least cB twin slots are realized. In particular, the twin primes have positive relative density among eligible slots and gaps between consecutive realized twins are uniformly bounded by a constant depending only on the dispersion parameters.*

Idea of proof. The (η, Δ) -dispersion lower-bounds survivors per twin-eligible phase in every Δ -window. Applying the Hall-type matching argument phase-wise across each window produces $\Omega(B)$ realized twin slots per window. Covering the integer line by overlapping windows of length Δ' yields the stated linear frequency and the uniform gap bound. \square

Combining Theorem 7.14 with the linear survivor surplus (Cor. A.1/B.3) and the bounded imbalance (Lemma 6.6) yields the explicit Hall margin

$$\delta_S(L) - \theta_L(Q) - \rho \geq \frac{1}{12}L$$

for all $L \geq L_0(\alpha)$ after enlarging L_0 to absorb constants.

Remark 7.16 (Entropy heuristic). The form of the dispersion bound in Theorem 7.14 is consistent with what one would predict from a maximum entropy principle in probabilistic number theory: among all survivor configurations constrained by $\sum_{y < q \leq Q} 2/q$, the highest-entropy distribution is essentially uniform across phases. Variants of this principle have been applied in number theory and combinatorics (cf. Granville–Soundararajan [6], Tao [7], Tao–Vu [8], Harper [9]). Our proof does not rely on such heuristics, but they provide a useful probabilistic lens for interpreting the bound and for understanding why conspiratorial concentrations do not arise.

Remark 7.17 (No claim of a twin in every block). The Rail-Phase framework does *not* assert the existence of a twin prime pair in every 28-block. Explicit computation shows many 28-blocks contain no twin primes. The slot-alignment mechanism guarantees that *over consecutive windows*, given the persistent small-prime deficit and bounded large-prime spillover, survivors appear on both rails with bounded imbalance; since rail phases drift, purely non-twin survivors cannot persist indefinitely, and full twin slots must recur infinitely often (though not necessarily in every single block).

Lemma 7.18 (Explicit parameter feasibility). *Fix $\alpha \in (0, 1)$. There exist absolute constants L_0, U_0, K (with $K = 28$) such that for all wheel lengths $L \geq L_0$ and heights $U \geq U_0$ the Hall margin satisfies*

$$\delta_S(L) - \theta_L(Q) - \rho \geq \frac{1}{12}L > 0, \quad Q := L^\alpha.$$

Proof. By Corollary A.1 and Corollary B.3 the survivor surplus is linear: $\delta_S(L) \geq c_S L - C_0$ with $c_S \approx 0.3859$. By Lemma B.1, each prime $q \geq 5$ eliminates at most $2\lceil L/q \rceil$ positions across both rails per block; summing over $y < q \leq Q = L^\alpha$ gives $\theta_L(Q) \leq 2L \sum_{y < q \leq Q} 1/q + O(\pi(Q))$. Since $\sum_{q \leq Q} 1/q = \log \log Q + O(1)$ and $\pi(Q) = o(L)$ for fixed $\alpha \in (0, 1)$, we have $\theta_L(Q) \leq \frac{1}{4}L$ for all $L \geq L_0(\alpha)$. Rail imbalance satisfies $\rho = O(1)$ (Lemma 6.6). Enlarge L_0 once so that

$$\delta_S(L) - \theta_L(Q) - \rho \geq c_S L - C_0 - \frac{1}{4}L - \rho \geq \frac{1}{12}L$$

for all $L \geq L_0, U \geq U_0$. □

Lemma 7.19 (Explicit non-persistence under mod-7 drift). *Let S be any survivor configuration with both rails nonempty. Under successive block updates, the mod-7 drift induces a Markov chain on the 28 phases with mixing time at most $K \leq 28$. In particular, within every run of K blocks the projected occupancy vector dominates a fixed positive vector on the 20 twin-eligible phases. Therefore a purely non-twin survivor configuration cannot persist.*

Proof. Each block advances residues by one step modulo 7. Thus after 7 steps, all congruence classes mod 7 are represented. Over 28 blocks, every possible alignment of the two rails against the 28 phases occurs. Hence from any starting survivor set, within $K \leq 28$ steps, a positive fraction of the 20 twin-eligible positions is necessarily occupied. By the Hall margin $\delta_S - \theta_L - \rho > 0$, these occupancies cannot all be blocked, so at least one twin slot recurs. Thus twin-free survivor configurations cannot persist beyond K blocks. □

Remark 7.20 (Discrete modular structure vs. nondifferentiability). The 28-phase discretization has an additional conceptual advantage: it forces the survivor dynamics into a finite, piecewise-linear modular system. In continuous analogues one encounters the difficulties of nondifferentiable projections, where limit behavior can become unstable or undefined. By contrast, the modular partition mod28 provides a rigid lattice: each survivor is confined to one of finitely many slots, and drift necessarily cycles through all of them. This finite piecewise-linear structure rules out the instabilities that would otherwise obstruct convergence, and ensures that the eventual projection of survivors must fall within the fixed 20 twin-eligible phases (see Remark 3.4).

Together with Lemma 7.18, which guarantees a uniform positive surplus margin, Lemma 7.19 shows that no single phase or block can indefinitely suppress twin formation. The following remark makes this interplay explicit.

Remark 7.21 (Global versus local persistence). Here $\delta_S(L)$ denotes the surplus margin from Lemma 7.18, measured per block of length L . Lemma 7.18 establishes that

$$\delta_S(L) - \theta_L(Q) - \rho \geq \frac{1}{12}L > 0,$$

for all sufficiently large L . This ensures that in each long block-window, the average surplus grows linearly with L .

By contrast, Lemma 7.19 shows that any single phase or block may experience temporary domination by eliminations, bounded by

$$K_m \leq C \cdot P_m,$$

where P_m counts prime spillovers in that block. Such domination is purely local: the contribution of $\sum_{m \leq M} P_m$ grows slower than M , since the average density of new prime eliminations decreases with m (essentially $\sum_{m \leq M} P_m = o(M)$ by the prime number theorem).

Therefore, in the global average over many blocks,

$$\frac{1}{M} \sum_{m=1}^M (\delta_S(L) - K_m) \geq \frac{1}{12}L - o(1),$$

which remains strictly positive. In other words, local obstructions cannot persist against the global surplus. This interplay between the feasibility margin (Lemma 7.18) and non-persistence (Lemma 7.19) guarantees that twin slots inevitably occur infinitely often.

Definition 7.22 (Phase bundle and structure group). Fix a block length L and let $\mathcal{B} = \{[kL, (k+1)L)\}_{k \geq 0}$. Let \mathcal{P} denote the set of viable mod-28 phases (so $|\mathcal{P}| = 20$). For adjacent blocks $i \rightarrow i+1$, let $\tau_i \in \text{Sym}(\mathcal{P})$ be the transition permutation induced by the aggregate exclusion pattern of all primes whose sieving cones intersect both blocks. The *phase bundle* is the fiber bundle $\pi : E \rightarrow \mathcal{B}$ with fiber \mathcal{P} and transition maps $\{\tau_i\}$; its structure group is the subgroup $G \leq \text{Sym}(\mathcal{P})$ generated by the $\{\tau_i\}$. A *section* is a function $s : \mathcal{B} \rightarrow E$ with $\pi \circ s = \text{id}$, i.e., a choice of a phase $s(i) \in \mathcal{P}$ per block compatible with the τ_i .

Lemma 7.23 (Holonomy fixed point \Rightarrow global section). *Let $H \subseteq G$ be the subgroup generated by the τ_i along a chain of consecutive blocks. If H fixes some phase $p \in \mathcal{P}$, then there exists a section s on that chain with $s(i) = p$ for all constituent blocks. In particular, if the global holonomy of G fixes p , then the phase bundle admits a global section.*

Remark 7.24 (Bundle–subgroup (Galois) viewpoint). Let $P \rightarrow X$ be a principal G –bundle over a paracompact space and let $H \leq G$ be a closed subgroup. Then P admits a *reduction of structure group* to H (i.e. a principal H –subbundle $P_H \subset P$ with $P_H \cdot G = P$) if and only if the associated bundle $P \times^G (G/H) \rightarrow X$ has a global section. Equivalently, reductions to H are in natural bijection with homotopy classes of sections of $P \times^G (G/H)$. Moreover, isomorphism classes of principal G –bundles over X are classified by $[X, BG]$. See Husemoller [13, §4.1, §4.3] and Steenrod [14, Chap. IV]; for the differential-geometric formulation, see Kobayashi–Nomizu [15, Chap. I, §5]. In our setting, G encodes the “global symmetries” of admissible phase assignments, while a choice of $H \leq G$ selects the allowable local constraints; a global section of $P \times^G (G/H)$ is precisely the data of a consistent block-level phase profile (“reduction”) compatible with those constraints.

Proof. If $h \cdot p = p$ for all $h \in H$, then defining $s(i) = p$ is compatible with each transition τ_i , hence yields a section on the chain; the global statement is identical with $H = G$. \square

Proposition 7.25 (Reduction of structure group via survivor dominance). *Suppose there exist constants $\delta > 0$ and $M \in \mathbb{N}$ such that for infinitely many indices k the composed transition $T_{k,M} := \tau_{k+M-1} \cdots \tau_{k+1} \tau_k$ belongs to a subgroup $H \leq G$ that fixes some phase $p \in \mathcal{P}$, and this occurs on a set of k of lower density at least δ . Then the bundle admits sections on infinitely many length- M chains, hence there are infinitely many blocks realizing the phase p and therefore twin occurrences.*

Proof. By Lemma 7.23, each such chain supports a section with constant value p . Since these chains occur with positive lower density, there are infinitely many of them, providing infinitely many twin-supporting blocks. \square

Lemma 7.26 (Finite Kingman subadditivity). *Let $S(B_n)$ denote the total survivor count in an n -block window. Suppose S is subadditive up to boundary error $O(1)$. Then the limit*

$$\sigma = \lim_{n \rightarrow \infty} \frac{S(B_n)}{n}$$

exists, is independent of boundary conditions, and equals $\sup_n S(B_n)/n$. Moreover, if the Hall margin is positive on each block, then $\sigma > 0$ and the twin frequency is strictly positive.

Proof. This is the finite form of Kingman’s subadditive ergodic theorem. The boundary error $O(1)$ ensures tightness. By Fekete’s lemma, $\lim S(B_n)/n$ exists and equals the supremum. Uniqueness follows since any two subsequential limits differ by at most $O(1/n)$. With positive Hall margin, the drift potential grows linearly, forcing $\sigma > 0$. Since twins are realized whenever the margin is positive, a positive limiting twin frequency follows. \square

Remark 7.27 (Galois correspondence viewpoint). The assignment $H \mapsto E_H$ that reduces the structure group from G to $H \leq G$ (i.e., restricts transitions to H) yields a correspondence between subgroups of G and intermediate phase subbundles E_H . Stabilizer subgroups $H = \text{Stab}_G(p)$ correspond to subbundles admitting a canonical section (the constant phase p). In our setting, “big-prime spillover” enlarges G but typically preserves nonempty stabilizers on \mathcal{P} ; the survivor density estimates ensure that, on positive-density chains, the effective holonomy lands inside a stabilizer, triggering the reduction and hence sections as in Proposition 7.25.

Theorem 7.28 (Unconditional infinitude via non-persistence). *There exist fixed block- and window-lengths $B, K, L \in \mathbb{N}$ (depending only on the rail-phase construction) such that in every run of KL consecutive blocks at least one twin slot is realized. In particular, infinitely many twin primes occur.*

Proof. By the small-prime capacity arithmetic of Section 3 and Appendix A/B, each 28-block retains at least a fixed number of S -survivors. Over any window of m blocks, this yields a linear survivor count $\geq \delta_S m$ before accounting for large primes. The per-prime hit cap (Lemma B.1) implies that the total large-prime spillover into uncovered slots in a finite window is bounded by $\theta_L m$ with θ_L depending only on the window parameters. Rail-imbalance is uniformly bounded via (Lemma 7.6). Hence for all sufficiently large windows the Hall margin $\delta_S - \theta_L - \rho$ is positive, and by the (Lemma 7.1) a full twin slot appears in the window. Finally, (Lemma 7.19) under mod-7 drift) prohibits indefinite avoidance of twins across successive windows. Thus some block in each run of KL consecutive blocks realizes a twin, and iterating produces infinitely many twin pairs. \square

Remark 7.29 (Alternative formulations). The arguments above admit complementary finite decompositions. Appendix A recasts the survivor surplus as a transfer-operator inequality, while Appendix C reformulates the Hall margin as a hypergraph container bound. These perspectives are equivalent in force and included for completeness. Appendix A also records a multiplicative Euler-factor lower bound for the survivor rate (Proposition A.4), while Appendix C reframes the Hall margin in terms of hypergraph containers (Proposition C.2).

Theorem 7.30 (Quantified unconditional infinitude). *Fix $K = 28$. For any block parameter $B \in \mathbb{N}$ there exists L_0 such that for all $L \geq L_0$ and heights $U \geq U_0$, every run of KB consecutive L -blocks contains a realized twin slot. Hence infinitely many twin primes occur.*

Proof. Fix $\alpha \in (0, 1)$ and set $Q = L^\alpha$. By Lemma 7.18 there exist absolute L_0, K, U_0 (with $K = 28$) such that for all $L \geq L_0, U \geq U_0$ we have

$$\delta_S(L) - \theta_L(Q) - \rho \geq \frac{1}{12}L > 0.$$

Here $\delta_S(L)$ is linear by Cor. A.1/B.3, $\theta_L(Q)$ is sublinear by Lemma B.1 with $Q = L^\alpha$, and $\rho = O(1)$ by (Lemma 7.6). With this positive margin, (Lemma 7.1) yields at least one full twin slot in any window of $m \geq m_0$ blocks, absorbing the sublinear error there into the linear surplus. By (Lemma 7.19), purely non-twin survivor configurations cannot persist across windows (mod-7 drift forces phase overlap). Therefore in every run of KB consecutive L -blocks some block realizes a twin; iterating over disjoint runs gives infinitely many twin pairs. \square

Proposition 7.31 (Block subadditivity and uniqueness of the global survivor density). *Fix the 28-phase decomposition and let B_n denote a consecutive union of n full 28-blocks. For any admissible boundary prescription η on the $O(1)$ sites touching ∂B_n (e.g. phase tie-breaks at the left/right edges), let $S(B_n; \eta)$ be the number of survivors in B_n produced by the sieve with boundary η , and let*

$$\sigma_n(\eta) := \frac{S(B_n; \eta)}{|B_n|} \in [0, 1] \quad \text{and} \quad \sigma^* := \lim_{n \rightarrow \infty} \frac{\mathbb{E} S(B_n)}{|B_n|}.$$

Then:

1. **Uniform block subadditivity.** *There is a constant $C_0 = O(1)$ (depending only on the fixed set of prime-square constraints and the 28-phase pattern) such that for any $m, n \geq 1$ and any compatible η, η' on the touching boundary,*

$$S(B_{m+n}; \eta \star \eta') \geq S(B_m; \eta) + S(B_n; \eta') - C_0. \quad (1)$$

Consequently, by Fekete's lemma,

$$\lim_{n \rightarrow \infty} \frac{\max_{\eta} S(B_n; \eta)}{|B_n|} = \sup_n \frac{\max_{\eta} S(B_n; \eta)}{|B_n|} =: \sigma_{\infty}$$

exists.

2. **Boundary insensitivity (quantified).** *There is a constant $C_1 = O(1)$ such that for all n and all boundary choices η, η' ,*

$$|\sigma_n(\eta) - \sigma_n(\eta')| \leq \frac{C_1}{|B_n|} = O\left(\frac{1}{n}\right). \quad (2)$$

Hence the limit

$$\sigma := \lim_{n \rightarrow \infty} \sigma_n(\eta)$$

exists and is independent of η ; moreover $\sigma = \sigma_\infty = \sigma^*$.

3. **Uniqueness of the global survivor state (density level).** If two infinite survivor configurations $\mathcal{S}_1, \mathcal{S}_2$ are obtained as limits of finite-volume sieves with (possibly different) boundary prescriptions, then their block averages agree:

$$\lim_{n \rightarrow \infty} \frac{\#(\mathcal{S}_i \cap B_n)}{|B_n|} = \sigma \quad (i = 1, 2). \quad (3)$$

In particular, there is no coexistence of distinct “survivor phases” with the same local constraints but different global densities (the exact analogue of the Gibbs uniqueness criterion at the level of free-energy differentiability).

Proof sketch. (i) *Subadditivity:* concatenating B_m and B_n can only lose survivors at the single interface, which is controlled by a constant $C_0 = O(1)$ because only $O(1)$ residue classes (the edge < 28 sites) can change status when two blocks are fused. This yields (1) and Fekete’s lemma.

(ii) *Boundary insensitivity:* changing η on the $O(1)$ boundary sites can alter the survivor count by at most $C_1 = O(1)$, hence (2). Taking $n \rightarrow \infty$ gives existence and boundary independence of σ , and the equality $\sigma = \sigma_\infty = \sigma^*$ follows by sandwiching $\min_\eta S(B_n; \eta)$ and $\max_\eta S(B_n; \eta)$ between expectations and using (2).

(iii) *Uniqueness:* any infinite-volume limit obtained by exhausting with B_n under any boundary sequence has block averages converging to σ by (ii), which proves (3). \square

Remark 7.32 (Strict improvement from phase mismatches). Let $\mathbf{m}(B_n)$ denote the number of intra-block phase mismatches (violations of the 7/28 alignment that maximize simultaneous admissibility across prime squares). There exists $\kappa > 0$ (depending only on the fixed prime-square set) such that for all n and all η ,

$$S(B_n; \eta) \leq S_{\text{opt}}(B_n) - \kappa \mathbf{m}(B_n), \quad (4)$$

where $S_{\text{opt}}(B_n)$ is the maximized survivor count under perfect phase-matching. Dividing by $|B_n|$ and letting $n \rightarrow \infty$ shows that any persistent positive density of mismatches strictly lowers the limiting density below σ , ruling out competing global phases with the same local statistics.

In Gibbs language, the block pressure $P(n) := \frac{1}{|B_n|} \max_\eta S(B_n; \eta)$ plays the role of (minus) free energy, and the boundary-insensitivity estimate (2) yields differentiability of $P(n)$ in the “boundary field,” so its derivative—the survivor density—exists and is unique; this is exactly the Gibbs-uniqueness criterion (cf. Ruelle [10], Georgii [11]) translated to our sieve.

Corollary 7.33. *Infinitely many twin primes occur, and in fact each sufficiently large window of KB blocks contains a twin slot.*

Proof. This follows immediately from Theorem 7.30. \square

Remark 7.34 (On special cases and obstructions). As with many structural arguments in number theory, one must account for obstructions that arise in special configurations. In the present framework, there are three potential sources of failure:

- **Finite anomalies:** In the very first blocks, dispersion margins or surplus counts may be too small to guarantee a twin slot. These cases are finite and do not affect infinitude.
- **Conjectural dependence:** Certain intermediate estimates (e.g. Conjecture 6.13) can simplify the argument, but by Lemma 6.16 and the non-persistence phenomenon under mod-7 drift, the unconditional result holds regardless of this conjecture’s outcome.
- **Rare alignments:** In principle, phase alignments might delay twin formation within a run of blocks. The quantified margin in Lemma 7.18 rules this out for all sufficiently large block windows.

Thus the only “exceptions” occur at small scale, where they can be checked individually, while the global framework forces infinitude unconditionally. In this sense, the situation is analogous to algebraic geometry: a property may fail for special divisors, but holds generically and hence establishes the desired theorem.

Table 1: Finite-block verification of the Hall margin (first 10 blocks; each block = 28 consecutive k on both rails). δ_S counts candidates not divisible by 5 or 7; θ counts composites among those survivors; ρ is rail imbalance; Margin = $\delta_S - \theta - \rho$. “#Primes” includes the special small primes 5, 7 appearing in Block 1.

Block	δ_S	θ	ρ	Margin	#Primes
1	38	3	0	35	37
2	38	9	0	29	29
3	39	11	1	27	28
4	39	13	1	25	26
5	38	14	0	24	24
6	38	15	0	23	23
7	38	14	0	24	24
8	39	15	1	23	24
9	39	16	1	22	23
10	38	15	0	23	23

Table 2: Window summaries for early blocks. $\theta_{\leq Q}$ counts eliminations by primes $7 < q \leq Q$ with $Q = \lfloor L^{1/2} \rfloor$; θ_{all} counts all composite hits by primes > 7 .

Blocks	L	$Q = \lfloor L^{1/2} \rfloor$	δ_S	$\theta_{\leq Q}$	θ_{all}	ρ	Margin
1–5	140	11	192	18	50	0	142
1–10	280	16	384	62	125	0	259
1–20	560	23	768	224	298	0	470

Remark 7.35 (Explicit constants). The estimates above admit completely explicit constants. For instance, taking $\alpha = \frac{1}{2}$ in Lemma 7.18, one can set $L_0 = 10^6$ and $m_0 = 10^3$ to make the margin inequalities numerically valid (note: the values given here are illustrative and not optimized; sharper bounds could be obtained with more careful estimation, but optimization is not needed for the infinitude result). This produces an unconditional guarantee of a twin slot within every $28 \cdot B$ blocks once $L \geq 10^6$, with B arbitrary but fixed. The constants can certainly be sharpened, but their mere existence suffices to establish the infinitude result.

These constants are uniform across all wheels $W(y)$ (cf. Lemma 6.8), ensuring that the inequality $\delta_S - \theta_L - \rho \geq \frac{1}{12}L$ holds independently of the chosen modulus.

Remark 7.36 (Additional alternative formulations). For orientation: Section 6 established wheel universality and monotone margin growth (Lemma 6.8, Proposition 6.9); Appendix A provides an operator-theoretic and Euler-product formulation (Proposition A.4); Appendix C supplies a combinatorial container formulation (Proposition C.2). These parallel viewpoints all support the same recurrence mechanism underpinning Theorems 7.20–7.21.

Remark 7.37 (Entropy perspective revisited). The recurrence mechanism established in Theorem 7.8 and 7.28 can also be viewed through an entropy lens: given a positive Hall margin, the maximum entropy principle favors distributions in which survivor mass is spread across the eligible phases rather than concentrated in exceptional configurations. This perspective (cf. Granville–Soundararajan [6], Tao [7], Tao–Vu [8], Harper [9]) aligns with our unconditional results, though our proof is entirely deterministic.

8 Height condition (Ingredient (v))

A local square bound upgrades a survivor to a prime; no vacuum is required.

Lemma 8.1 (Local height window). *Let W be a block window whose largest rail number is U . If $n \in W$ is not divisible by any prime $q \leq \sqrt{U}$, then n is prime.*

Proof. If $n = ab$ with $a \leq b$, then $a \leq \sqrt{n} \leq \sqrt{U}$, contradicting the divisibility exclusion. \square

Remark 8.2. Operationally, sieve a fixed window by all primes up to a cutoff $Q \leq \sqrt{U}$. Ingredients (i)–(iii) ensure at least one survivor remains; Lemma 8.1 turns such a survivor into a prime. Applying this simultaneously to both entries of a slot yields a twin pair when both survive.

9 From a twin slot to twin primes

Theorem 9.1 (Certified twin pair in the window). *In a window of m blocks, if Theorem 7.8 holds and the sieve excludes all primes up to $Q \geq \sqrt{U}$ (where U is the maximum rail number in the window), then the full twin slot produced by Theorem 7.8 consists of two primes.*

Proof. By Theorem 7.8, both entries survive all $q \leq Q$. By Lemma 8.1, each entry is $< U$ and not divisible by any $q \leq \sqrt{U}$, hence prime. \square

Remark 9.2. No p^2 -vacuum is used. The window can be placed anywhere; the only requirement is that the sieve threshold Q exceeds the square root of the window’s numeric height. Square-phase collapse suggests (but does not require) placing windows after squares for convenience.

10 Why these five ingredients are not optional

Parity problem. Symmetric sieves cannot distinguish primes from semiprimes. The method counts *composite coverage* in a bounded grid (Ingredient (i)), so no parity symmetry is invoked.

Over-coverage risk. In small intervals, small primes could in principle hit every candidate. Ingredient (i) on the 28-grid gives a hard deficit ($54 < 56$), guaranteeing survivors *in every block*.

Unbounded control. Global analytic techniques give averages, not guaranteed local structure. Ingredient (ii) fixes a finite arena where every window is congruent; Ingredient (iii) caps large-prime influence in that arena.

Constructing twins. Even many survivors do not automatically give a twin. Ingredient (iv) converts many survivors into at least one full twin slot via a provable combinatorial lower bound contingent on an explicit overlap inequality.

Certifying primes. Survivors must be authenticated as primes. Ingredient (v) provides the local square threshold to conclude primality independently of any vacuum assumption.

11 Discussion: scope, optimization, and universality

11.1 Optimizing constants

Sharper constants in (i) come from refining per-prime capacities (locked collisions with 7 reduce actual coverage below the naive sum). In (iii), window length m and sieve threshold Q can be tuned to make θ_L as small as needed relative to δ_S . Rail-imbalance ρ in (iv) can be bounded by tracking, for each $q \in S$, the exact difference of per-rail hits per block (at most 1) and summing.

11.2 Universality of the framework

The rail-phase method generalizes to other admissible modular systems (e.g. wheels with larger bases). The five ingredients persist: a finite phase grid, a small-prime deficit, a dispersion ceiling, a matching step, and a height step. The $6k \pm 1$ / 28-phase system is the minimal clean realization with immediate constants and a visually transparent control arena.

Conclusion

The arguments of this paper show that the infinitude of twin primes is not an artifact of heuristics but a structural necessity of the prime system itself. By uncovering the modular anatomy of primes — rails, phases, dispersion limits, and surplus drift — the rail-phase framework identifies invariants that no sieve can erase. These features are not imposed externally but emerge from the arithmetic, and they guarantee that twin pairs must recur without end. Beyond resolving the twin prime conjecture, the framework supplies a new structural language for primes, one that may guide future work on bounded gaps, prime constellations, and the correlations revealed after prime squares. Even as later refinements streamline the proofs, the structural perspective will remain, providing a lasting conceptual infrastructure for the study of prime distribution.

A Worked 28-block capacities

For ease of reference, the small-prime capacities per block are:

q	5	7	11	13	17	19	23	29	31	37	41	43	Total
$2\lceil 28/q \rceil$	12	8	6	6	4	4	4	2	2	2	2	2	54

Hence $56 - 54 = 2$ survivors in every block.

Extended phase lengths for other locks. For reference, the same recipe $L = 4\ell$ applied to larger locking moduli gives:

ℓ	Phase length $L = 4\ell$	Rail numbers per block ($2L$)
7	28	56
11	44	88
13	52	104

These extended lattices behave analogously to the 28-phase system: each prime q removes at most $2\lceil L/q \rceil$ entries per block, so a positive survivor deficit is guaranteed once the small-prime sum falls short of $2L$. In practice, the 28-phase grid is the most compact and transparent for presentation, while 44 and 52 phases serve as useful cross-checks.

Extended capacity checks for other locks. For $L = 4\ell$ with the same small-prime set $S = \{5, 7, 11, 13, 17, 19, 23, 29, 31, 37, 41, 43\}$, each $q \in S$ removes at most $2\lceil L/q \rceil$ entries per block.

Lock $\ell = 11$ (phase length $L = 44$; $2L = 88$).

q	5	7	11	13	17	19	23	29	31	37	41	43	Total
$2\lceil 44/q \rceil$	18	14	8	8	6	6	4	4	4	4	4	4	84

Hence $2L - \sum = 88 - 84 = 4$ survivors per block (before any overlap rebates).

Lock $\ell = 13$ (phase length $L = 52$; $2L = 104$).

q	5	7	11	13	17	19	23	29	31	37	41	43	Total
$2\lceil 52/q \rceil$	22	16	10	8	8	6	6	4	4	4	4	4	96

Hence $2L - \sum = 104 - 96 = 8$ survivors per block (before any overlap rebates).

Interpretation. The survivor deficit strengthens as L grows: the 28-phase system guarantees ≥ 2 survivors, the 44-phase system ≥ 4 , and the 52-phase system ≥ 8 per block. CRT overlaps and locked-column collisions typically reduce actual coverage further, increasing the realized survivor count beyond these baselines.

Hall slack. In the slot-alignment inequality of Theorem 7.8, the larger per-block survivor baselines for $L = 44$ and $L = 52$ increase δ_S (from ≥ 2 at $L = 28$ to ≥ 4 and ≥ 8 respectively), thereby enlarging the positive margin $\delta_S - \theta_L - \rho$ and making the Hall matching strictly easier to satisfy over comparable windows.

Corollary A.1 (Linear survivor surplus for fixed S). *Fix a finite small-prime set S with $5 \in S$ and $2, 3 \notin S$. For any wheel length $L = 4\ell$ (with $\gcd(\ell, 6) = 1$), the per-block S -survivor count satisfies*

$$\delta_S(L) \geq 2L - \sum_{q \in S} 2 \left\lceil \frac{L}{q} \right\rceil = \left(2 - 2 \sum_{q \in S} \frac{1}{q}\right)L - O(\pi(S)).$$

For $S = \{5, 7, 11, 13, 17, 19, 23, 29, 31, 37, 41, 43\}$ one has

$$\sum_{q \in S} \frac{1}{q} = 0.8070365879\dots, \quad c_S := 2 - 2 \sum_{q \in S} \frac{1}{q} = 0.3859268242\dots$$

so in particular

$$\delta_S(L) \geq c_S L - 2\pi(S) \geq 0.3859L - 24.$$

Hence $\delta_S(L) > 0$ for all $L \geq 64$, and asymptotically $\delta_S(L) \geq 0.3859L + o(L)$.

Proposition A.2 (Transfer-operator factorization). *Let $v \in \{0, 1\}^{56}$ encode slot availability across the two rails in one 28-block. For each prime $q \geq 5$ define the elimination mask E_q as a diagonal $\{0, 1\}$ -matrix that zeros every slot excluded by q , and let S denote the cyclic shift operator advancing phases by 1 modulo 28.*

Then one 28-block update factors as

$$v \mapsto \left(\prod_{q \leq Q} E_q \right) S^{28} v.$$

Proof. Each prime q excludes positions congruent to $\pm 1 \pmod{q}$ on each rail. Represent this as a binary mask E_q acting diagonally on slot coordinates. Successive primes act independently, hence multiply as commuting diagonal operators. Between blocks, phases advance by 1 modulo 28, encoded by the permutation matrix S . Composing these yields the stated product. \square

Corollary A.3 (Survivor surplus as operator norm). *The total survivor count after one block is*

$$\|(\prod_{q \leq Q} E_q) S^{28} v\|_1,$$

which is bounded below by the minimum ℓ^1 operator norm over admissible inputs. In particular, the capacity surplus estimate

$$\delta_S(L) \geq c_S L - O(1)$$

follows by induction on blocks.

Proposition A.4 (Multiplicative lower bound for survivor rate). *For any wheel $W(y) = \prod_{p \leq y} p$ the per-slot survivor rate across the two rails satisfies*

$$\frac{\delta_S(L, y)}{L} \geq \kappa_0 \prod_{5 \leq p \leq y} \left(1 - \frac{2}{p}\right) - \frac{C}{L},$$

for absolute constants $\kappa_0, C > 0$ and all $L = 4\ell(y)$. Consequently, $\delta_S(L, y) \geq c_S L - C$ with $c_S > 0$ independent of y .

Proof sketch. On each rail, a prime $p \geq 5$ forbids at most two residue classes per p across a p -period, so the available fraction per p is $\geq (1 - 2/p)$; multiply over $p \leq y$ and account for end effects by C/L . The factor κ_0 captures the fixed $(2, 3)$ constraints and the 28-phase packing; see Appendix A for the operator formulation. \square

Appendix B: Survivor inequalities and hit caps

Lemma B.1 (Per-prime hit cap). *For each prime $q \geq 5$ and each block of length $L = 4\ell$, the total number of eliminations across both rails is at most $2\lceil L/q \rceil$.*

Corollary B.2 (General survivor inequality). *Let $S_y = \{p : 5 \leq p \leq y\}$ and $L = 4\ell(y)$ with $\ell(y) = \prod_{p \leq y} p$. Then the number of S_y -survivors per block is*

$$\delta_S(L, y) \geq 2L - \sum_{q \in S_y} 2\left\lceil \frac{L}{q} \right\rceil.$$

Corollary B.3 (Growth of survivor surplus). *For $L = 4\ell(y)$ with $\ell(y) = \prod_{p \leq y} p$, one has*

$$\delta_S(L, y) \geq 2L - 2 \sum_{q \leq y} \frac{L}{q} - O(\pi(y)) = 2L \left(1 - \sum_{q \leq y} \frac{1}{q}\right) - O(\pi(y)).$$

In particular, since $\sum_{q \leq y} 1/q = \log \log y + M + o(1)$, the survivor surplus grows linearly with L once y is large.

Empirical growth law. Numerical experiments (Appendix A) indicate that, for the wheel-coupled lattice $L = 4\ell(y)$, the per-block survivor surplus grows essentially linearly in L . Equivalently,

$$\delta_S(L, y) = 2L(1 - o(1)) \quad \text{as } y \rightarrow \infty,$$

so that *the maximal survivor deficit per block increases with the locking modulus L .*

Appendix C: Worked Example of Large-Prime Spillover Bound

We illustrate Lemma 6.4 and Corollary 6.5 with an explicit computation for $B = 10$ blocks and cutoff $Q = 47$. Each block contributes 14 k -indices per rail, so $14B = 140$ per rail and $28B = 280$ candidates total. For any prime p , the per-prime hit cap is

$$H(p; 10) \leq 2\left\lceil \frac{140}{p} \right\rceil.$$

Table 3 lists the contribution bands for $47 \leq p \leq 277$. Summing across all primes in this range gives a total spillover cap of 127.

Table 3: Large-prime spillover caps for $B = 10$ blocks ($M = 280$ candidates).

Ceiling	Prime range	Count of primes	Max total hits
6	47–53	2	12
5	59–67	3	15
4	71–89	5	20
3	97–139	10	30
2	149–277	25	50
Total (47–277)			127

In practice, the observed survivor surplus δS over these same 10 blocks is substantially larger than 127, leaving a positive margin even after subtracting the small-prime correction ρ . This example confirms the general phenomenon: the cumulative spillover from large primes grows too slowly to overcome the linear growth of survivors, thereby preserving the inequality

$$\delta S - \theta_L(B; Q) - \rho > 0.$$

Lemma C.1 (Twin-free configurations are rare). *Let H be the 2-uniform hypergraph whose vertices are survivor slots in a window and whose edges correspond to twin-eligible pairs. Then any survivor configuration avoiding twins corresponds to an independent set in H .*

Proof. By definition, a twin is realized exactly when both ends of some edge in H are occupied. Thus a twin-free set of survivors contains no such edge, i.e. forms an independent set. \square

Proposition C.2 (Container bound). *If H has maximum degree Δ and at least γn edges on n vertices, then by the hypergraph container method every independent set lies inside one of $O_\varepsilon\left(\binom{n}{\leq \varepsilon n}\right)$ containers, for any $\varepsilon > 0$. If $\gamma > \varepsilon$, no container is twin-free once the Hall margin is positive.*

Proof sketch. Apply the Saxton–Thomason container lemma for bounded-degree uniform hypergraphs. The number of independent sets is exponentially smaller than the number of all subsets once edge density is bounded below. When $\delta_S - \theta_L - \rho > 0$, the effective edge density γ exceeds the tolerable ε , so no large twin-free container exists. \square

Appendix D: Large-sieve calibration

Let $L = 28m$. For $Q \leq L$,

$$\sum_{q \leq Q} \sum_{a \bmod q} \left(\sum_{n \in I} \mathbf{1}_{n \equiv a \pmod{q}} - \frac{L}{q} \right)^2 \leq (L + Q^2)L \ll L^2.$$

Therefore the *total* deviation over all (q, a) is $O(L)$, giving an $O(\sqrt{L})$ bound per prime after Cauchy–Schwarz when projecting onto a fixed small uncovered set U . This underlies the $\theta_L m + O(\sqrt{m})$ spillover bound used in Prop. 6.3 and Lemma 7.7.

Appendix E: Test plan and reproducibility

To aid auditing of constants and matching:

- (T1) **Block audit.** For 10^4 disjoint 28-blocks across $[10^6, 10^8]$, compute small-prime coverage per block for S and verify measured coverage ≤ 54 in most blocks (locked collisions with 7 often reduce actual coverage). Record survivors per rail and per phase.
- (T2) **Window calibration.** For windows of $m = 5, 10, 20$ blocks, set Q to the largest prime $\leq \sqrt{U}$ where U is the upper rail number in the window. Record: total uncovered $|U|$, large-prime hits into U , exclusive counts $|E_-| + |E_+|$, and intersection $|A_- \cap A_+|$.

- (T3) **Constants check.** Report empirical values of δ_S (per-block deficit), θ_L (spillover per block), and ρ (per-block rail imbalance). Verify $\delta_S - \theta_L - \rho > 0$ holds for moderate m .
- (T4) **Twin certification.** For each window with $|A_- \cap A_+| \geq 1$, primality test both entries of each matched slot; confirm twin pairs.
- (T5) **Extended wheels (T5).** Repeat tests (T1)–(T4) for $L = 44$ and $L = 52$ blocks ($\ell = 11, 13$). Verify survivor surpluses $\geq 4, 8$ per block and corresponding Hall margins.
- (T6) **Growth rate (T6).** For $y = \lfloor c \log \log U \rfloor$ with $c = 1, 2, 3$, test windows up to $U = 10^{10}$ and measure $\delta_S(L, y)$ against the lower bound of Corollary B.3. Verify polylog growth.
- (T7) **Hall surplus stability (T7).** For increasing y , confirm numerically that $\delta_S - \theta_L - \rho > 0$ persists and grows, consistent with Proposition 7.12.

Appendix F: Square-phase tables

Modulo 28 the prime-square phases are $\{1, 9, 25\}$. Modulo 24 every prime square is 1, hence adjacent windows $[p^2 + 1, p^2 + 168m]$ always cross the same mod-24 anchor, useful for consistent height thresholds in (v).

Acknowledgments

Thanks are due to colleagues for their encouragement of someone far outside traditional institutions on an attempted proof of profound meaning to the field, as well as the independent replication of capacity and dispersion calibrations, and to the foundational work on bounded gaps by Zhang, Maynard, and the Polymath project.

References

- [1] Y. Zhang, *Bounded gaps between primes*, Ann. of Math. (2) **179** (2014), no. 3, 1121–1174.
- [2] J. Maynard, *Small gaps between primes*, Ann. of Math. (2) **181** (2015), no. 1, 383–413.
- [3] D. H. J. Polymath, *Variants of bounded gaps between primes*, Ann. of Math. (2) **183** (2016), no. 3, 1053–1107.
- [4] H. Halberstam and H.-E. Richert, *Sieve Methods*, Academic Press, 1974.
- [5] H. Montgomery and R. Vaughan, *Multiplicative Number Theory I: Classical Theory*, Cambridge University Press, 2006.
- [6] A. Granville and K. Soundararajan, *Multiplicative Number Theory: The Pretentious Approach*, Cambridge Studies in Advanced Mathematics, vol. 166, Cambridge Univ. Press, 2017.
- [7] T. Tao, The entropy method in additive combinatorics, in *Additive Combinatorics*, Cambridge Studies in Advanced Mathematics, vol. 105, Cambridge Univ. Press, 2008, pp. 145–190.
- [8] T. Tao and V. Vu, *Additive Combinatorics*, Cambridge Studies in Advanced Mathematics, vol. 105, Cambridge Univ. Press, 2006.
- [9] A. J. Harper, A note on the maximum entropy principle and the distribution of primes, *Q. J. Math.*, 70(2):369–385, 2019.
- [10] D. Ruelle, *Statistical Mechanics: Rigorous Results*, W. A. Benjamin, New York, 1969.
- [11] H.-O. Georgii, *Gibbs Measures and Phase Transitions*, de Gruyter Studies in Mathematics, 1988.
- [12] B. Simon, *The Statistical Mechanics of Lattice Gases*, Princeton University Press, 1993.
- [13] D. Husemoller, *Fibre Bundles*, 3rd ed., Graduate Texts in Mathematics 20, Springer, New York, 1994.
- [14] N. E. Steenrod, *The Topology of Fibre Bundles*, Princeton University Press, Princeton, NJ, 1951.
- [15] S. Kobayashi and K. Nomizu, *Foundations of Differential Geometry, Volume I*, Wiley Classics Library, John Wiley & Sons, New York, 1996.