

Research Article

Prime Curvature Geometry and the Structure of Additive Prime Deviations

Bill Riemers¹

1. Independent researcher

This paper develops a geometric framework for analyzing remainder terms in additive prime and prime-like problems. The approach is based on a multiplicative decomposition of the deviation into an explicit small-prime residue structure and a medium-to-large-prime tail exhibiting intrinsic exponential curvature. This decomposition leads to the formulation of the Prime Curvature Geometry Hypothesis (PCGH), together with an associated curvature constant Ω , intended to govern pointwise remainder envelopes beyond the scope of the classical Fundamental Lemma of sieve theory.^[1] Specializing the framework to the Goldbach problem yields the Prime Curvature Geometry Conjecture for Goldbach (PCGC–Goldbach) and an explicit curvature constant Ω_{prime} . The resulting conjectural bounds are compatible with Hardy–Littlewood A predictions in the asymptotic limit, while providing a geometric mechanism that organizes deviation behaviour across finite analytic windows and cutoff scales.

An explicit bounding envelope for the Goldbach remainder is derived and shown to control the exact remainder uniformly. This envelope is proved not to be an asymptotic proxy: the exact remainder exhibits non-vanishing exponential curvature within cutoff cells, preventing pointwise asymptotic equivalence. This distinction clarifies the limits of scale-based approximations and isolates the geometric source of deviation.

The present paper is restricted to the derivation and formal statement of the geometric framework, associated hypotheses, and bounding envelopes. Numerical validation, certified bounds over large finite ranges, and reduction theorems relating PCC–Goldbach to Hardy–Littlewood–A and Bombieri–Vinogradov–type hypotheses^{[2][3]} are developed in subsequent work.

Correspondence: papers@team.qeios.com — Qeios will forward to the authors

1. Introduction

1.1. Motivation for a General Geometric Framework

Additive problems involving primes share a common structural feature: their main terms are governed by local congruence conditions imposed by the small primes, while their error terms reflect global fluctuations in the distribution of primes. Classical conjectures such as the Hardy–Littlewood prime–pair conjectures incorporate this observation by introducing finite correction factors derived from the sieve.

However, the Hardy–Littlewood framework is not a universal predictive principle. Its successes and limitations both indicate that a deeper, more structural description of analytic remainders is required.

HL–A Encodes Small-Prime Structure But Does Not Generalize

For problems such as Goldbach representations, the normalized Hardy–Littlewood predictions match empirical data with striking accuracy. This agreement is explained by the fact that Harding–Littlewood Conjecture A (HL–A) captures precisely the finite congruence restrictions enforced by the small primes.

Yet HL–A fails to organize the full behaviour of primes. It does not supply a coherent global model of remainder terms; it ties the abundance of twin primes to that of “lucky primes”; and it provides no mechanism for describing fluctuations across multiple analytic scales. Thus HL–A is an effective *particular solution* for certain additive problems, but it is not a *general framework* from which all such predictions can be derived.

This limitation is a primary motivation for first developing a general geometric model, and only afterwards tuning it to the Goldbach setting.

Primorial Structure and Scale Separation

The sieve of Eratosthenes imposes deterministic exclusion patterns modulo the primorials $p_k^{\#}$. As k increases, these patterns refine but still arise from the same finite rule set. This hierarchy induces a natural separation of scales in the behavior of primes:

- **Small primes:** rigid congruence obstructions.
- **Intermediate scales:** structured but non-deterministic fluctuations, responsible for the “curvature” phenomena observed in localized prime correlations.
- **Large scales:** weak correlations resembling statistically independent noise.

Any analytic model that respects sieve structure must therefore separate remainder terms according to these scales.

The Need For a Remainder Geometry.

Empirical data from Goldbach pair counts, prime tuples, and short intervals show that analytic deviations cannot be represented by a single global error term. Rather, they decompose into

- a *small-prime factor* capturing deterministic congruence structure,
- a *curvature term* describing intermediate-scale deviations,
- and a *tail term* reflecting large-scale fluctuations.

HL–A implicitly models only the first of these components. A general “remainder geometry” is required to describe the other two.

Why the General Model Must Precede the Goldbach Case

If one begins directly with the Goldbach problem, any proposed correction or curvature equation risks appearing ad hoc. In contrast, if the deviation structure is derived from a general geometric decomposition compatible with sieve constraints, then the Goldbach geometry becomes the *unique specialization* of the framework. HL–A appears as one particular class of main-term corrections, and the resulting conjecture becomes structurally motivated and falsifiable.

In this sense, HL–A motivates the general geometric model precisely because it fails to generalize: its partial successes identify the small-prime contribution that any model must contain, while its limitations reveal the necessity of a broader framework.

1.2. Goals of This Paper

The objectives of this work are:

- to develop a hypothesis of general curvature geometry of primes from first principles,
- to specialize this framework to the Goldbach setting,
- to state PCGC–Goldbach,
- to provide the theoretical foundation for empirical validation and reduction theorems (treated in subsequent papers).

1.3. Main Statements

Goldbach's conjecture asserts that every sufficiently large even integer $2n$ can be expressed as a sum of two primes. Let

$$G(2n) = \sum_{m \in \mathbb{Z}} 1_{\text{prime}}(n-m) 1_{\text{prime}}(n+m)$$

denote the number of such Goldbach representations. The Hardy–Littlewood A conjecture (HL–A) [4] predicts the asymptotic behaviour

$$G(2n) \sim \mathfrak{S}_{GB}(2n) \frac{2n}{\log^2(n)},$$

where $\mathfrak{S}_{GB}(2n) := 2C_2 \mathfrak{S}(2n)$ is the twin prime constant C_2 multiplied by the classical singular series.

Despite overwhelming heuristic support, particularly after incorporating Hardy–Littlewood circle-method corrections, HL–A remains unproved and lies beyond the reach of classical sieve methods [1]. The obstruction is the well-known sieve barrier: after sieving by primes up to $x^{1/2}$, no further cancellation is available to control prime correlations at the scale required for HL–A. Nevertheless, extensive heuristic and computational evidence indicates that Hardy–Littlewood–type predictions accurately describe Goldbach pair counts across wide finite ranges; see, for example, the computational validation in [5]. This tension between strong empirical agreement and the limitations of existing analytic tools motivates the search for additional underlying structure.

This barrier suggests that any successful approach must exploit additional structure beyond classical sieve estimates.

The purpose of the present paper is not to prove HL–A, but to derive a *geometric framework* that governs deviations from Hardy–Littlewood–type predictions in additive prime problems. Goldbach's problem serves as a motivating example, but the framework developed here applies more generally to prime-like sets and structured configuration families.

Then, for a symmetric window of radius M , define the localized Goldbach count

$$G(2n; M) = \sum_{0 < |m| \leq M} 1_{\text{prime}}(n-m) 1_{\text{prime}}(n+m),$$

and let $\mathring{G}(2n; M)$ denote a corresponding Hardy–Littlewood–type predictor restricted to the same window.

The deviation

$$\varepsilon(2n, M) := G(2n; M) - \mathring{G}(2n; M)$$

may be viewed as a function on the two-parameter space $(2n, M)$.

New contribution. This paper introduces a *prime curvature geometry* for such deviations. The central hypothesis is that remainder terms in additive prime problems admit a multiplicative decomposition separating:

- a finite structural factor determined by very small primes, and
- a tail factor arising from medium and large primes, whose dominant behaviour is governed by a renormalized curvature constant.

This decomposition induces a natural geometric interpretation of remainder growth across analytic scales and explains the emergence of exponential-type bounds in short-interval problems.

The goal of this paper is to derive this geometric framework and formulate the corresponding general curvature hypothesis. Its specialization to Goldbach, numerical validation, and reduction theorems are treated in after this paper.

2. Fundamental Definitions

2.1. Euler–Cap

The Euler–cap constraint was introduced and motivated in previous work ^[5] as a structural regulator for finite–scale Goldbach analysis. In the present work it is adopted uniformly as an admissibility condition on window parameters; its use here is inherited from that framework rather than introduced for the purposes of the present conjecture.

Definition 1 (Euler–Cap and Admissible Window Size).

For a given central value n , define the *Euler cap* by

$$Ecap(n) := \frac{(2n+1) - \sqrt{8n+1}}{2n}.$$

The Euler cap determines the maximal fractional size of a symmetric window about n for which the mapping

$$m \mapsto (n-m)(n+m)$$

remains injective, with both factors positive and distinct.

Given a window parameter $\alpha \in (0, 1]$, the associated window radius is defined by

$$M := \lfloor \min(\alpha, Ecap(n))n \rfloor.$$

Thus $Ecap(n)$ enforces a structural upper bound on admissible window sizes, while α allows additional restriction when required.

Unless explicitly stated otherwise, all window parameters in this paper are assumed to satisfy the Euler–cap constraint. Since $Ecap(n) \rightarrow 1$ as $n \rightarrow \infty$, the choice $\alpha = 1$ is asymptotically admissible and may be used safely in large- n regimes.

The Euler–cap restriction is therefore treated as part of the ambient analytic framework throughout this paper, and not as a conjectural or optimizable parameter. Its effect of bypassing small-prime weight sensitivity is a convenient but secondary consequence.

More generally, this serves as an example of the normalization issues that must be addressed when applying the PCGH framework to the formulation of new conjectures.

2.2. Geometric Terms

Many variables in this paper are treated as *geometric parameters* rather than as functions in the analytic sense. Just as one would not normally write $x = r(x, y, z)\cos\theta(x, y, z)$, we avoid functional notation for quantities that are best understood as coordinates, lengths, or changes of scale.

Our primary coordinate system is (n, m) . The variable n denotes a position along the central axis, while m denotes a displacement from that position. Capital letters such as M and N represent *lengths* measured along the same axis, rather than coordinates themselves.

The parameter

$$L := \sqrt{2M}$$

is introduced as a change of scale corresponding to the effective length of a short interval. It should not be interpreted as a function of M , but rather as a coordinate transformation to a different scale. Accordingly, L is treated as a continuous parameter.

When it is necessary to relate the window size to the central coordinate, a dimensionless scale parameter α is introduced. This parameter specifies a window length via

$$M = n\alpha,$$

and should be interpreted as a relative scale, not as a slope or direction in the (n, m) -plane. In particular, α converts a coordinate magnitude into a length; it does not define a line or trajectory in the coordinate

space.

In contexts where α varies with n , $\alpha(n)$ denotes a varying scale, so that $M = n\alpha(n)$. This represents a family of window sizes indexed by n , rather than a curve in the (n, m) -plane.

Because this framework mixes discrete variables with continuous scale parameters, care is required when interpreting non-integer quantities. Unless stated otherwise, an expression of the form $n\alpha(n)$ is interpreted as $\lfloor n\alpha(n) \rfloor$ when discretization is required and the Euler–cap is not in effect, and otherwise as specified in Definition 1. By contrast, the parameter L is treated as continuous; applying a floor operation to L generally leads to inaccurate results and is avoided.

2.3. Admissible Parity

The parity restriction was introduced in previous work [5]. Here it is used purely as an admissibility convention for windowed sums and products.

Definition 2 (Parity–Admissible Window Index Set).

Let $n \in \mathbb{N}$ and let $M \geq 1$. The symmetric window is defined as

$$I_M := \{m \in \mathbb{Z} : 0 < |m| \leq M\}.$$

The *parity–admissible index set* is defined as

$$I^{par}(n; M) := \{m \in I_M : n + m \equiv 1 \pmod{2}\}.$$

Equivalently, $I^{par}(n; M)$ consists of those shifts m with $0 < |m| \leq M$ for which both $n - m$ and $n + m$ are odd. Unless stated otherwise, all summations over window variables m in this paper are implicitly restricted to $m \in I^{par}(n; M)$.

2.4. Goldbach Singular Series Factors

Definition 3 (Goldbach Singular Series Factors).

For an even integer $2n$, the *local semiprime correction factor* [4] is defined as

$$\mathfrak{S}(2n) := \prod_{\substack{p \in \mathbb{P} \setminus \{2\} \\ p \mid n}} \frac{p-1}{p-2}.$$

The *prime–pair constant* [4] is defined as

$$C_2 := \prod_{p \in \mathbb{P} \setminus \{2\}} \left(1 - \frac{1}{(p-1)^2}\right).$$

The corresponding *Goldbach singular series factor* [4] is defined as

$$\mathfrak{S}_{GB}(2n) := 2C_2 \mathfrak{S}(2n).$$

2.5. Complementary and Full Euler–Type Products

In addition to the local semiprime correction factor $\mathfrak{S}(2n)$ defined above, it is occasionally convenient to refer to the complementary and full Euler–type products obtained by modifying the divisibility condition on the prime index.

Definition 4 (Complementary and Full Local Semiprime Products).

The *complementary* local semiprime product is defined as

$$\mathfrak{S}^C(2n) := \prod_{\substack{p \in \mathbb{P} \setminus \{2\} \\ p \nmid n}} \frac{p-1}{p-2},$$

and the corresponding *full* product is defined as

$$\mathfrak{S}^{\cdot}(2n) := \prod_{p \in \mathbb{P} \setminus \{2\}} \frac{p-1}{p-2}.$$

Convention (Notation). A superscript C indicates replacement of the divisibility condition $p \mid n$ by $p \nmid n$, while a superscript \cdot indicates removal of the divisibility condition entirely. No additional structure is implied.

2.6. Effective Local Modulus

Definition 5 (Effective Local Modulus).

Let $n \in \mathbb{N}$ and let p be an odd prime. Then the minimum contributing prime is defined as

$$p_{\min}(n) := \begin{cases} 3, & 3 \mid n, \\ 5, & 3 \nmid n. \end{cases}$$

Definition 6 (Admissible Primes for Effective Moduli).

For each $n \geq 2$, let

$$P_{\text{eff}}(n) := \{p \in \mathbb{P} : p \geq p_{\min}(n)\}.$$

Throughout this paper, $\mathcal{Q}_p(n)$ is regarded as defined only for $p \in P_{\text{eff}}(n)$.

For any odd prime q_{\min} , define the partial Euler product

$$\mathcal{Q}_p^{(q_{\min})} := \prod_{\substack{q \in \mathbb{P} \\ q_{\min} \leq q \leq p}} (q-1).$$

In analogy with the prime-indexed quantity $\mathcal{Q}_p(n)$, it is useful to name the cumulative small-prime contribution associated to a cutoff scale L . Accordingly, the envelope-scale effective local modulus is defined as

$$\mathcal{Q}_p(n) := \mathcal{Q}_p^{(p_{\min}(n))}.$$

Thus $\mathcal{Q}_p(n)$ encodes the cumulative residue structure imposed by the odd primes up to p , with the only n -dependence arising from the choice of base prime $p_{\min}(n)$ according to whether $3 \mid n$. This convention is tailored to the singular-series geometry underlying Goldbach-type problems.

For the geometric analysis of monotonic changes in aggregated quantities, it is useful to introduce a natural envelope scale along the n -axis. Let $\{p_i\}$ denote the increasing sequence of odd primes starting at $p_{\min}(n)$. Then the envelope length associated to the prime interval $[p_i, p_{i+1})$ is defined as

$$\begin{aligned} N_i^{(3)} &:= \left(\mathcal{Q}_{p_{i+1}}^{(3)} \right)^2 - \left(\mathcal{Q}_{p_i}^{(3)} \right)^2, \\ N_i^{(5)} &:= \left(\mathcal{Q}_{p_{i+1}}^{(5)} \right)^2 - \left(\mathcal{Q}_{p_i}^{(5)} \right)^2. \end{aligned}$$

These quantities represent the natural block sizes over which the local residue structure induced by the primes below p_{i+1} remains fixed.

The squaring reflects the Goldbach geometry $(n-m)(n+m) = n^2 - m^2$, so that changes in the effective modulus $\mathcal{Q}_p(n)$ induce corresponding blocks along the n -axis. Within each such envelope, aggregated quantities are expected to exhibit coherent (and typically monotonic) behavior.

For convenience, and for use in short-interval and envelope-scale expressions, we define the effective local modulus at scale L by

$$\mathcal{Q}(2n; L) := \prod_{\substack{p \in \mathbb{P}_{\text{eff}}(n) \\ \mathcal{Q}_p(n) \leq L}} (p-1) = \prod_{\substack{p \in \mathbb{P}_{\text{eff}}(n) \\ \mathcal{Q}_p(n) \leq L}} (p-1).$$

Convention (Interpretation of Cutoff Inequalities). Throughout this paper, whenever an inequality or cutoff condition of the form $\mathcal{Q}_p(n) \leq L$ is used, it is understood that the effective local modulus $\mathcal{Q}_p(n)$ is defined with the base prime $p_{\min}(n)$ fixed as in Definition 5. No restriction on p is imposed beyond $p \geq p_{\min}(n)$.

2.7. Additional Series Operators

The classical Hardy–Littlewood singular series $\mathfrak{S}(2n)$ is naturally an asymptotic object. For finite–window analysis it is convenient to introduce auxiliary operators that partition the series into base and tail

components relative to a square cutoff.

Definition 7 (Cutoff Operators).

Define the square cutoff operator

$$\mathfrak{S}q(x) := x^2.$$

The function $\mathcal{Q}(\cdot)$ may also be used as a cutoff operator.

We use the tags *head* and *tail* to denote the base (small-prime) and tail (large-prime) components relative to a given cutoff.

Definition 8 (Cutoff Components of \mathfrak{S}).

Let $M > 0$ and $n \geq 2$. Then define

$$\begin{aligned}\mathfrak{S}_{\text{head}}^{\text{Sq}}(2n; M) &:= \prod_{\substack{p \in \mathbb{P} \setminus \{2\} \\ p \mid n \\ p^2 \leq M}} \frac{p-1}{p-2}, \quad \mathfrak{S}_{\text{tail}}^{\text{Sq}}(2n; M) := \prod_{\substack{p \in \mathbb{P} \setminus \{2\} \\ p \mid n \\ p^2 > M}} \frac{p-1}{p-2}. \\ \mathfrak{S}_{\text{head}}^{\mathcal{Q}}(2n; L) &:= \prod_{\substack{p \in \mathbb{P}_{\text{eff}}(n) \\ p \mid n \\ \mathcal{Q}_p(n) \leq L}} \frac{p-1}{p-2}, \quad \mathfrak{S}_{\text{tail}}^{\mathcal{Q}}(2n; L) := \prod_{\substack{p \in \mathbb{P}_{\text{eff}}(n) \\ p \mid n \\ \mathcal{Q}_p(n) > L}} \frac{p-1}{p-2}. \\ \mathfrak{S}^{\text{Sq}}(2n) &:= \mathfrak{S}(2n), \\ \mathfrak{S}^{\mathcal{Q}}(2n) &:= \prod_{\substack{p \in \mathbb{P}_{\text{eff}}(n) \\ p \mid n}} \frac{p-1}{p-2}.\end{aligned}$$

Empty products are interpreted as 1.

Identity 1 (\mathfrak{S} Head–Tail Products).

$$\begin{aligned}\mathfrak{S}(2n) &= \mathfrak{S}^{\text{Sq}}(2n) = \mathfrak{S}_{\text{head}}^{\text{Sq}}(2n; M)\mathfrak{S}_{\text{tail}}^{\text{Sq}}(2n; M) \quad \forall M > 0, \\ \mathfrak{S}^{\mathcal{Q}}(2n) &= \mathfrak{S}_{\text{head}}^{\mathcal{Q}}(2n; L)\mathfrak{S}_{\text{tail}}^{\mathcal{Q}}(2n; L) \quad \forall L > 0.\end{aligned}$$

Proof. These identities follow by design of having the *tail* term continue immediately after the *head* term ends. \square

Identity 2 (\mathfrak{S} Equality).

$$\mathfrak{S}^{\text{Sq}}(2n) = \mathfrak{S}^{\mathcal{Q}}(2n)$$

Proof. Both products are taken over primes dividing n , with admissibility restricted to $p \in \mathbb{P}_{\text{eff}}(n)$ as in Definition 18. The only term that could appear in $\mathfrak{S}^{\text{Sq}}(2n)$ but not in $\mathfrak{S}^{\mathcal{Q}}(2n)$ is $p = 3$.

However, $\mathfrak{S}^{\text{Sq}}(2n)$ contains a $p = 3$ factor only if $3 \mid n$. In that case $p_{\min}(n) = 3$, so $3 \in \mathbb{P}_{\text{eff}}(n)$, and the corresponding term also appears in $\mathfrak{S}^{\mathcal{Q}}(2n)$. Thus both products range over the same prime divisors of n ,

and are identical. \square

Definition 9 (Effective Moduli Interval Max).

$$\begin{aligned}\mathcal{Q}(2n; L) &:= \mathcal{Q}_{P_0(2n; L)}(n), \\ P_0(2n; L) &:= \max \{p \in P_{\text{eff}}(n) : p \mid n, \mathcal{Q}_p(n) \leq L\}, \\ &\text{with } P_0(2n; L) := p_{\min}(n) \text{ if the set is empty.}\end{aligned}$$

Identity 3 (Equivalent Tail Terms).

$$\mathfrak{S}_{\text{tail}}^{\mathcal{Q}}(2n; L) = \mathfrak{S}_{\text{tail}}^{\text{Sq}}(2n; (P_0(2n; L))^2) \quad \forall L \geq p_{\min}(n) - 1.$$

Proof. Both tail products are taken over primes $p \in P_{\text{eff}}(n)$ dividing n . By definition of $P_0(2n; L)$, the condition $\mathcal{Q}_p(n) > L$ is equivalent to $p > P_0(2n; L)$. Hence both sides are products over the same index set

$$\{p \in P_{\text{eff}}(n) : p \mid n, p > 2, p > P_0(2n; L)\},$$

and therefore agree. \square

Identity 4 (Equivalent Header Terms).

$$\mathfrak{S}_{\text{head}}^{\mathcal{Q}}(2n; L) = \mathfrak{S}_{\text{head}}^{\text{Sq}}(2n; (P_0(2n; L))^2) \quad \forall L \geq p_{\min}(n) - 1.$$

Proof. Both header products are taken over primes $p \in P_{\text{eff}}(n)$ dividing n . By definition of $P_0(2n; L)$, the condition $\mathcal{Q}_p(n) \leq L$ is equivalent to $p \leq P_0(2n; L)$. Hence both sides are products over the same index set

$$\{p \in P_{\text{eff}}(n) : p \mid n, p > 2, p \leq P_0(2n; L)\},$$

and therefore agree. \square

The above identities do not necessarily hold for \cdot and \mathcal{C} operators. As the \mathcal{Q} operator will retain exclusion of primes not in $P_{\text{eff}}(n)$, while the Sq operator will essentially flip an exclusion to an inclusion. As such the following identity is added:

Identity 5 (\mathfrak{S} complement conversion at the p_{\min} cutoff).

For every $n \geq 2$,

$$\mathfrak{S}^{\text{Sq}, \mathcal{C}}(2n) = \mathfrak{S}_{\text{head}}^{\text{Sq}, \mathcal{C}}(2n; (p_{\min}(n) - 1)^2) \mathfrak{S}^{\mathcal{Q}, \mathcal{C}}(2n) \quad \forall n \geq p_{\min}(n).$$

2.8. HL–A Predictor Corrected for Density and Short Intervals

Definition 10 (Short–Interval HL–A Predictor).

For Goldbach configurations $(n - m, n + m)$ with local window $0 < |m| \leq M$, define the Hardy–Littlewood Conjecture A predictor (including the standard circle–method correction) (HL–Windowed) for *ordered*

pairs by

$$\hat{G}^{HL}(2n; M) := 2C_2 \mathfrak{S}(2n) \sum_{m \in I^{par}(2n; M)} \omega(n-m)\omega(n+m).$$

The decomposition of \mathfrak{S} into base and tail components and the resulting restriction of prime contributions are specific to the present work. The use of weighted summation in place of the classical $x/\log^2 x$ density factor was introduced in previous work [5] and is retained here to permit finite-scale bounds.

3. A General Predictive Geometry for Prime–Like Sets

Many problems in analytic number theory can be formulated as follows: one seeks to count how often certain structured candidate configurations fall into specified “prime–like” sets, in order to construct a suitable analytic prediction for that count and to understand the deviation between measurement and prediction. This section establishes a unified framework for that purpose and serves as the general backbone for the specialized Goldbach geometry in later sections.

3.1. Prime–Like Sets and Candidate Configurations

Let

$$S_1, S_2, S_3, \dots$$

be a sequence of *prime–like sets*.

Definition 11 (Prime–Like Set).

A subset $S_j \subset \mathbb{N}$ is called *prime–like* if membership is determined by a multiplicative divisibility–type rule; that is, there exists a $\{0, 1\}$ –valued indicator $\mathcal{D}_j(n)$, depending only on residue constraints for n , such that

$$n \in S_j \iff \mathcal{D}_j(n) = 1.$$

Examples include: primes, primes in fixed congruence classes, integers avoiding a set of forbidden prime factors, survivors of iterated sieves, or any structured set defined by local multiplicative restrictions.

Let

$$C = \{C_i : i \in I\}$$

be a family of *candidate vectors*, where

$$C_i = \left(C_i^{(1)}, C_i^{(2)}, C_i^{(3)}, \dots \right).$$

Each coordinate $C_i^{(j)}$ is to be tested for membership in the corresponding prime–like set S_j .

3.2. Measured Joint Incidence

The total measured incidence is defined by

$$G(C; S_1, S_2, \dots) := \sum_{i \in I} \sum_{j \geq 1} \mathbf{1}_{S_j}(C_i^{(j)}),$$

which counts how many components of how many candidate vectors lie in the respective prime–like sets.

This formulation includes, as special cases:

- Goldbach configurations $(n - m, n + m)$ with $S_1 = S_2 = \{\text{primes}\}$,
- twin–prime patterns $(n, n + 2)$,
- prime k –tuples,
- survivors of restricted sieves,
- cross–correlation patterns of primes across multiple offsets.

3.3. The Prediction Problem

The objective is to construct a closed–form analytic predictor

$$\hat{G}(C; S_1, S_2, \dots),$$

which satisfies:

- dependence only on the structural residue data defining the sets S_j ;
- multiplicative behavior analogous to Hardy–Littlewood type predictions;
- incorporation of both small–prime constraints (CRT–type) and medium–prime geometric fluctuations;
- accurate specialization to the case where the S_j are genuine primes.

This is the *general prediction problem*, of which the classical Hardy–Littlewood heuristics are a special instance.

3.4. Deviation

Once \hat{G} is defined, the deviation is set as

$$\varepsilon(C; S_1, S_2, \dots) := G(C; S_1, S_2, \dots) - \hat{G}(C; S_1, S_2, \dots).$$

Two questions then arise:

1. How large can ε be, relative to the structure of the sets S_j ?
2. How does the geometry of the small and medium primes dictate the size, curvature, and oscillatory behaviour of these deviations?

The framework developed in the sequel interprets deviation not as an analytic error, but as a *geometric curvature quantity* built multiplicatively from contributions of the residue structure of the small primes and a tail factor encoding medium-prime fluctuations.

3.5. Purpose and Scope of This Section

The remainder of this section develops:

- a CRT-inspired multiplicative base product capturing small-prime structure across the sets S_j ;
- a medium-prime tail factor producing a universal curvature term;
- a corresponding curvature constant governing asymptotic fluctuations;
- a generalized Hardy–Littlewood–type prediction functional \hat{G} ;
- a deviation bound of the form

$$|\varepsilon| \leq (\text{window normalization}) \times (\text{curvature term}).$$

This provides the structural foundation for the specialized Goldbach geometry studied later and prepares the ground for reductions and applications in subsequent sections.

3.6. Predicted Counts

Given a family of candidate vectors $C = \{C_i\}$ and prime-like sets S_1, S_2, \dots ; consider the task of constructing an analytic prediction for the joint incidence count

$$\hat{G}(C; S_1, S_2, \dots).$$

Two classical approaches exist, one based on local density heuristics and one based on sieve theoretic product series, but neither is adequate on its own. A blended multiplicative–geometric predictor will ultimately emerge as the appropriate generalization.

3.6.1. Density-Based Prediction

Using a method specified by Iwaniec and Kowalski, if the coordinates of the candidates $C_i^{(j)}$ behave with little correlation, one may attempt a purely density-based estimate [1]:

$$\hat{G}_{\text{density}}(C; S_1, S_2, \dots) := \sum_i \prod_j \omega(C_i^{(j)}, S_j),$$

where $\omega(x, S_j)$ denotes the local density or “probability weight” for an integer x to lie in the prime-like set S_j . Such estimates work well for uncorrelated random models, but even mild arithmetic correlations can render them unreliable. The multiplicative structure $\prod_j \omega(C_i^{(j)}, S_j)$ implicitly assumes independence across coordinates, but arithmetic progressions modulo small primes create systematic dependencies that violate this assumption. These correlations accumulate multiplicatively across the product, leading to systematic over- or under-estimation that cannot be corrected by normalization alone. Correct normalization of the weights ω is also essential.

3.6.2. Pointwise Sieve Prediction

A second approach mentioned in works by Harman and by Iwaniec and Kowalski uses pointwise sieve estimates [6][1], typically expressed as product series describing local sieving constraints:

$$\hat{G}_{\text{sieve}}(C; S_1, S_2, \dots) := \sum_i \prod_j \left(1 - f(C_i^{(j)}, S_1, S_2, \dots)\right) P_j(S_j).$$

Here $P_j(S_j)$ encodes the local exclusion probability induced by the small primes defining S_j , while the correction factor $f(\cdot)$ accounts for residual terms, local obstructions, and normalization discrepancies arising in the sieve. pointwise sieves capture local structure accurately, but they do not correctly handle medium-prime correlations or global normalization.

3.6.3. A Combined Predictor

A more accurate predictive functional blends density weighting with explicit sieve corrections [6][1], yielding

$$\hat{G}(C; S_1, S_2, \dots) := \sum_i \prod_j \left(1 - f(C_i^{(j)}, S_1, S_2, \dots)\right) P_j(S_j) \omega(C_i^{(j)}, S_j).$$

This form simultaneously

- incorporates local sieving structure through the multiplicative correction factors $1 - fP_j$,

- adjusts global normalization via the density weights ω ,
- and captures both rigid small-prime obstructions and coarse medium-prime statistical behavior.

In the special case where the sets S_j correspond to the primes, this functional reduces to a generalized Hardy–Littlewood–type predictor.

For such an expression to be meaningful beyond heuristic averaging, the resulting prediction must be *asymptotically convergent* and *uniformly controlled* with respect to the underlying scale parameters. In particular, convergence must hold at the level of finite windows or logical groupings, rather than only after global averaging.

Remark. Failure of uniform asymptotic control at finite scales is precisely the obstruction encountered in classical sieve methods [11], and motivates the introduction of a geometric correction framework.

The failure of existing predictors to achieve uniform asymptotic control at finite scales motivates the introduction of a geometric correction, developed in the sections that follow.

3.7. Remainder Deficits

Given a candidate configuration family C and prime–like sets S_1, S_2, \dots , the *remainder deficit* is defined as

$$\varepsilon(C; S_1, S_2, \dots) := G(C; S_1, S_2, \dots) - \hat{G}(C; S_1, S_2, \dots).$$

There are two fundamentally different types of bounds one may attempt to place on ε :

(1) Statistical bounds.

Once G and \hat{G} are well understood on average, one can typically show that ε is statistically small. For Goldbach pairs, for instance, numerical data up to the limits of modern computations indicate that

$$|\varepsilon(2n; M)| = O(\sqrt{2M})$$

for the ranges of n that can be measured directly. Such bounds, while empirically robust, are extremely difficult to establish analytically.

(2) Uniform analytic bounds.

More interesting, and far more delicate, are *strict analytic* bounds valid for all sufficiently large C :

$$-R^-(C; S_1, S_2, \dots) \leq \varepsilon(C; S_1, S_2, \dots) \leq R^+(C; S_1, S_2, \dots).$$

By defining

$$R(C; S_1, S_2, \dots) := \max \left(R^-(C; S_1, S_2, \dots), R^+(C; S_1, S_2, \dots) \right),$$

one obtains the symmetric form

$$R(C; S_1, S_2, \dots) \geq |\varepsilon(C; S_1, S_2, \dots)|.$$

Because both G and \hat{G} decompose naturally over local blocks, the inequality can be refined pointwise:

$$R(C; S_1, S_2, \dots) \geq \left| \sum_i \varepsilon(C^{(i)}; S_1, S_2, \dots) \right|.$$

This immediately yields the trivial bound

$$R_{\text{trivial}}(C; S_1, S_2, \dots) \geq \sum_i \left| \varepsilon(C^{(i)}; S_1, S_2, \dots) \right|,$$

which is sharp only when all local deficits have the same sign. In general this overestimates the true remainder by double-counting correlated fluctuations between nearby blocks.

To obtain meaningful bounds, the remainder functional must be expressed in a form parallel to the predictor:

$$\hat{G}(C; S_1, S_2, \dots) := \sum_i \prod_j \left(1 - \mathcal{F}(C_i^{(j)}, S_1, S_2, \dots) P_j(S_j) \right) \omega(C_i^{(j)}, S_j).$$

When R shares the same structural decomposition, substantial portions of each term can be factored outside the summation, preventing the double-counting that plagues R_{trivial} .

At first glance, one might expect that the error in such expressions is governed by the Chinese Remainder Theorem (CRT), as summarized by Hardy and Wright [7]. However, direct CRT application to even simple problems such as Goldbach immediately predicts remainders of the order $O(n)$, which is entirely incompatible with observed behaviour. One remedy is to assume that all remainder contributions arise from a short interval. Using an interval of size $O(\sqrt{2M})$ indeed yields empirically correct bounds on the extremal error fluctuations, but only after introducing two additional elements:

- an independent tail factor accounting for primes not controlled by the short interval, and
- a measured curvature constant determining the effective interval scale.

While a single measured constant together with a new short-interval hypothesis *would* be sufficient for many applications, the aim here is a more universal and structurally transparent theory.

The remainder is therefore modelled through a *geometric* decomposition. The total remainder is written as a sum of local contributions:

$$R(C; S_1, S_2, \dots) := \sum_i R\left(C^{(i)}; S_1, S_2, \dots\right).$$

Each block $C^{(i)}$ is assigned a scalar parameter for the scale $L_i = \Lambda(C^{(i)})$, and the local remainder is postulated to factor as

$$R\left(C^{(i)}; S_1, S_2, \dots\right) = H(L_i; S_1, S_2, \dots)H_2(L_i; S_1, S_2, \dots)T(L_i; S_1, S_2, \dots)T_2(L_i; S_1, S_2, \dots).$$

The terms of Equation (62) will be defined in the following sections.

Remark (Which Set Does p Range Over?). Throughout this section the Euler products are written schematically, without fixing a single ambient set from which the primes p are drawn. This is intentional.

In the simplest situations, such as the Goldbach problem treated in later sections, there is a single prime–like set S , and all Euler products are taken over $p \in S$. In that case no ambiguity arises.

In more general configurations, however, different components of a candidate representation may be governed by *different* prime–like sets. For example, in Chen–type problems one variable is restricted to primes while the other is restricted to semiprimes, and the relevant sieving constraints are naturally expressed using more than one underlying set. In such cases the index set for p in each Euler product is the set appropriate to the factor being modelled: small–prime products H, H_2 range over primes enforcing local congruence restrictions, while the tail factors T, T_2 range over primes (or prime–like elements) whose influence is only felt through global density and curvature.

Formally, one may regard each Euler product as taken over an *effective prime–like support* determined by the candidate family C and the structural role of the factor under consideration. The precise specification of this support is part of the modelling choice and must be fixed on a case–by–case basis.

For clarity, all later specializations in this paper, and in particular the Goldbach specialization, use a single prime–like set, so that the indexing of the Euler products is unambiguous.

3.8. Construction of the Factors H, H_2, T , and T_2

The multiplicative decomposition (repeated from (62)

$$R\left(C^{(i)}; S_1, S_2, \dots\right) = H(L_i; S_1, S_2, \dots)H_2(L_i; S_1, S_2, \dots)T(L_i; S_1, S_2, \dots)T_2(L_i; S_1, S_2, \dots)$$

was motivated above as the only structure capable of avoiding the double-counting inherent in $R_{trivial} = \sum_i |\epsilon(C^{(i)})|$. A general method for constructing these factors is now outlined. This construction is not intended to be unique, but rather to isolate the distinct regimes of behaviour that naturally arise in sieve-type problems.

3.8.1. Small-Prime Structure: The Factors H and H_2

The small primes play a qualitatively different role from the medium and large primes. Their contribution to the remainder is not governed by a smooth Euler tail, but instead arises from *discrete residue effects*, parity constraints, and local structural obstructions. Unlike medium- and large-prime contributions, these effects do not average out over a local block $C^{(i)}$ and therefore require explicit treatment.

Crucially, the influence of the small primes decomposes into two distinct components:

- a *global*, residue-averaged contribution that is stable under summation and exponentiation, and
- a *local*, residue-sensitive correction that depends on the specific configuration of the block.

These are captured respectively by the factors H and H_2 .

For each configuration block $C^{(i)}$ with associated scale $L_i = \Lambda(C^{(i)})$, choose a cutoff

$$P_0 = P_0(L_i),$$

representing the largest prime whose residue behaviour remains structurally significant on the scale of $C^{(i)}$. Primes $p \leq P_0$ contribute through explicit small-prime structure, while primes $p > P_0$ are absorbed into the tail factors T and T_2 .

As before, assume that to each prime p and family of prime-like sets (S_j) a scale parameter is associated

$$\beta(p; S_1, S_2, \dots) > 0,$$

such that the small-prime regime for the block $C^{(i)}$ is defined by

$$p \leq P_0(L_i) \iff \beta(p; S_1, S_2, \dots) \leq L_i.$$

This compatibility between the cutoff $P_0(L_i)$ and the scale function β is essential: it ensures that all residue effects that fail to average on the scale L_i are treated explicitly, while all remaining prime contributions may be incorporated into the curvature tail.

$H(L_i)$ is a finite Euler product over the small primes $p \leq P_0(L_i)$ capturing the *global* small-prime sieving error. It incorporates no pointwise residue corrections, except for effects that have global structural impact (for example, those splitting the problem into a fixed finite family of cases independent of $P_0(L_i)$). The *global* small-prime contribution is defined by the finite Euler product

$$H(L_i; S_1, S_2, \dots) := \prod_{p \leq P_0(L_i)} h(p; S_1, S_2, \dots),$$

where $h(p; \cdot)$ is a residue-averaged multiplier describing the net sieving distortion introduced by the prime p . The factor H contains no pointwise residue information except for effects that have global structural impact (for example, splitting the problem into a fixed finite set of cases independent of P_0). At fixed scale L_i , H behaves as a constant and may be factored outside local summations.

$H_2(L_i)$ is a finite Euler product over the same primes $p \leq P_0(L_i)$ encoding *pointwise and residue-sensitive* structural effects arising from the small primes. The *local* small-prime correction is captured by

$$H_2(L_i; S_1, S_2, \dots) := \prod_{p \leq P_0(L_i)} h_2(p; L_i, S_1, S_2, \dots).$$

where $h_2(p; \cdot)$ encodes residue-sensitive and configuration-dependent effects of the prime p . Unlike H , this factor does not exponentiate cleanly and must remain attached to individual blocks. Its role is to correct the global constant H for local residue patterns that do not average out on the scale L_i .

Together, the factors H and H_2 represent the deficits of from small primes for our more general framework. The decomposition

$$R\left(C^{(i)}; S_1, S_2, \dots\right) = H(L_i)H_2(L_i)T(L_i)T_2(L_i)$$

correctly separates:

- global small-prime structure (H),
- local residue corrections (H_2),
- multiplicative curvature from medium primes (T), and
- higher-order and genuinely large-prime effects (T_2).

If the small-prime contribution were treated as a single factor, these roles would be conflated, leading to incorrect scaling and spurious overestimates of the remainder.

Remark. In the classical Hardy–Littlewood A formulation of Goldbach, the small-prime structure separates into a global constant C_2 and a residue-dependent correction $\mathfrak{S}_{sem}(2n)$. The present

decomposition HH_2 is directly analogous: H plays the role of the constant, while H_2 is the correct term to that constant. The advantage of the present framework is that this separation arises naturally from the geometric scale L of each block, rather than being imposed *a priori*.

3.8.2. The Medium–Prime Tail T (Unified Euler Product)

In classical sieve decompositions, small–prime effects are often separated into a finite factor, with the remaining primes treated as an exponential tail. For the present framework this separation is undesirable: it obscures the dependence on the cutoff scale and prevents a clean analytic description of the remainder.

Instead, the sieve correction is treated as a single Euler–type object with a scale–dependent cutoff. Let $P_0 = P_0(L_i)$ denote the cutoff prime associated with the configuration block $C^{(i)}$, typically determined by its scale parameter $L_i = \Lambda(C^{(i)})$. Primes $p \leq P_0$ contribute through explicit residue and parity structure, while primes $p > P_0$ contribute multiplicatively through a medium–prime tail.

The tail factor is defined by the Euler product

$$T(L_i; S_1, S_2, \dots) := \prod_{p > P_0(L_i)} t(p; S_1, S_2, \dots)^{L_i/\beta(p; S_1, S_2, \dots)},$$

where $t(p; \cdot)$ encodes the local contribution of the prime p .

To normalize this infinite product, introduce the global curvature constant

$$\Omega := \prod_p t(p; S_1, S_2, \dots)^{1/\beta(p; S_1, S_2, \dots)}.$$

The finite renormalization induced by the cutoff is then

$$\Theta(L_i; S_1, S_2, \dots) := \prod_{p \leq P_0(L_i)} t(p; S_1, S_2, \dots)^{-1/\beta(p; S_1, S_2, \dots)}.$$

With these definitions, the tail admits the normalized representation

$$T(L_i) = (\Omega \Theta(L_i))^{L_i},$$

where Ω captures the global prime–like structure, while $\Theta(L_i)$ accounts for the slowly varying correction associated with the finite cutoff. Formally, $\Theta(L_i) \rightarrow \Omega^{-1}$ as $P_0(L_i) \rightarrow \infty$.

Two key structural features follow immediately:

1. The quantity raised to the power L_i is the renormalized curvature constant. Define

$$\tilde{\Omega}(P_0) := \Omega \Theta(P_0), \Theta(P_0) := \prod_{p \leq P_0} t(p; S_1, S_2, \dots)^{-1/\beta(p; S_1, S_2, \dots)}.$$

All medium- and large-prime contributions appear only through this single renormalized constant $\tilde{\Omega}(P_0)$, while Ω remains the global curvature constant attached to the full infinite Euler product. No small primes are absorbed into T ; their effects are entirely contained in A .

2. The cutoff dependence of $\tilde{\Omega}(P_0)$ is structural rather than perturbative. Each time the cutoff $P_0(L_i)$ crosses a prime p , the factor $t(p)^{-1}$ moves from the infinite tail into the finite renormalization product.

The bracket

$$\tilde{\Omega}(L_i) := \Omega \prod_{p \leq P_0(L_i)} t(p; S_1, S_2, \dots)^{-1/\beta(p; S_1, S_2, \dots)}$$

therefore changes by an explicitly controlled multiplicative step at each cutoff transition. This is not a small perturbation in general; rather, it is the correct normalization ensuring that the truncated Euler product remains consistent with the underlying geometry. Because all cutoff-dependence is confined to the finite product, the tail retains the clean exponential form

$$T(L_i) = \tilde{\Omega}(L_i)^{L_i},$$

and the geometry remains stable under cutoff variation.

This unified formulation is essential. Because no small primes are artificially moved into H , the tail factor retains a coherent multiplicative structure, and the resulting curvature term $\tilde{\Omega}(P_0)$ becomes the natural geometric object governing remainder deviations.

In later sections, when specific sieve problems are introduced, this general tail structure will specialize to explicit formulas for $t(p)$, $P_0(L_i)$, and the resulting renormalized curvature constant.

3.8.3. The Medium and Large-Prime Correction T_2

Large primes (those exceeding both $P_0(L_i)$ and a second stability threshold) contribute only residue-level fluctuations. Their individual contributions are far below the natural scale of the problem, and their aggregate effect is conjecturally bounded.

A final term $T_2(L_i)$ is introduced with the following three axioms:

1. $T_2(L_i)$ is uniformly bounded above and below:

$$0 < c_1 \leq T_2(L_i) \leq c_2 < \infty,$$

2. $T_2(L_i)$ absorbs all higher-order correlations not handled by H , H_2 , or T ,
3. and in many applications, including the Goldbach case studied in Section 4, T_2 can be absorbed into the T term once an appropriate scaling is chosen.

Thus, $T_2(L_i)$ collects higher-order residue effects and genuinely large-prime corrections and conjecturally bounded:

$$T_2(L_i) := \prod_{p > P_0(L_i)} t_2(p; L_i, S_1, S_2, \dots)^{L_i/\beta(p; S_1, S_2, \dots)}.$$

3.8.4. Summary and General Hypothesis

The multiplicative decomposition

$$R(C^{(i)}) = H(L_i)H_2(L_i)T(L_i)T_2(L_i)$$

encapsulates, in a structured and modular way, the four principal sources of sieve deviation:

1. *Small primes.* $H(L_i)$ records the dominant residue-class distortions arising from primes p whose structural scale satisfies $\beta(p; S_1, S_2, \dots) \leq L_i$.
2. *Small primes corrections.* $H_2(L_i)$ records the dominant residue-class corrections arising from primes p whose structural scale satisfies $\beta(p; S_1, S_2, \dots) \leq L_i$. These are the primes for which local obstruction terms do not average out over the block $C^{(i)}$.
3. *Medium primes.* $T(L_i)$ collects the multiplicative tail of all primes with $\beta(p) > L_i$, and its leading behaviour is governed by a *renormalized curvature constant*

$$\tilde{\Omega}(P_0) := \Omega \Theta(P_0), \Theta(P_0) := \prod_{\beta(p) \leq L_i} t(p; S_1, S_2, \dots)^{-1/\beta(p; S_1, S_2, \dots)}.$$

Here Ω is the global curvature constant associated with the full (idealized) infinite product, while $\Theta(P_0)$ removes exactly the portion already absorbed into $A(L_i)$. The combined constant $\tilde{\Omega}(P_0)$ is the quantity that enters the exponent, and it remains stable under cutoff variation because $P_0 = P_0(L_i)$ grows only polylogarithmically in L_i .

4. *Medium and large primes and higher correlations.* $T_2(L_i)$ absorbs all remaining fluctuations, including large-prime residue effects and higher-order interactions, and is conjecturally bounded or slowly varying. It does not influence the exponential scale of $R(C^{(i)})$.

These components motivate the following hypothesis:

Hypothesis 1 (Prime Curvature Geometry Hypothesis (PCGH)).

Let S_1, S_2, \dots be any admissible family of prime-like sets, and let C be a structured configuration set admitting a decomposition into local blocks $C^{(i)}$, each equipped with a scale $L_i = \Lambda(C^{(i)})$.

Let $\hat{G}(C; S_1, S_2, \dots)$ be an admissible prediction functional satisfying the following conditions:

- (*Density normalization*) \hat{G} agrees with the expected cardinality of C under an uncorrelated prime-like model.
- (*Explicit small-prime structure*) \hat{G} incorporates all deterministic local sieve obstructions arising from primes p whose residue behaviour is structurally significant on the scale L_i .
- (*Finite-scale evaluability*) \hat{G} admits a well-defined evaluation on each block $C^{(i)}$, uniformly in L_i , without reliance on global averaging.

The analytic remainder is defined as

$$\varepsilon(C) := G(C) - \hat{G}(C).$$

Then $\varepsilon(C)$ admits a uniform bound of the form

$$|\varepsilon(C)| \leq R(C) := \sum_i H(L_i) H_2(L_i) T(L_i) T_2(L_i),$$

where the factors satisfy the following structural properties:

- (*Global small-prime factor*) $H(L_i)$ is a finite Euler product over primes p satisfying $\beta(p; S_1, S_2, \dots) \leq L_i$. It is residue-averaged, non-pointwise, and depends only on the scale L_i and the prime-like structure (S_j) .
- (*Local small-prime correction*) $H_2(L_i)$ is a finite product over the same primes, encoding residue-sensitive and configuration-dependent corrections that do not average out on the scale L_i .
- (*Geometric curvature tail*) $T(L_i)$ is a non-pointwise geometric term with exponential scaling

$T(L_i) = \left(\tilde{\Omega}(P_0(L_i)) \right)^{L_i}$, where $\tilde{\Omega}(P_0)$ is a finite curvature factor obtained by aggregating the contribution of primes $p > P_0(L_i)$. The role of $T(L_i)$ is to capture the cumulative geometric effect of the medium-prime tail at scale L_i .

- (*Residual corrections*) $T_2(L_i)$ consists of higher-order residue effects and genuinely large-prime contributions and remains uniformly bounded as $L_i \rightarrow \infty$.

Moreover, the decomposition above is *structural*: collapsing the small-prime factors H and H_2 into a single term destroys the scale separation required for a uniform bound.

Remark. No assumption is made in the Prime Curvature Geometry Hypothesis that the finite curvature factor $\tilde{\Omega}(P_0)$ converges to a nontrivial limit as $P_0 \rightarrow \infty$. In practice, for all physically meaningful choices of the local tail functions $t(p)$, one expects $\tilde{\Omega}(P_0) \rightarrow 1$, reflecting the rapid stabilization of large-prime contributions. Nonconvergent behavior would require deliberately nonphysical tail models and does not arise in standard prime-like settings.

Remark. The admissibility conditions above exclude predictors whose apparent accuracy arises solely from global averaging or implicit absorption of remainder terms. This restriction is essential for PCGH to address the classical sieve barrier.

Remark. The scale parameter $L_i = \Lambda(C^{(i)})$ is not restricted to contiguous intervals or uniform windows. It may represent *any logical grouping of candidates* appropriate to the problem at hand, including arithmetic progressions, divisor-class partitions, primorial blocks, or other structured subsets. This flexibility allows the Prime Curvature Geometry Hypothesis to be applied in settings where specialized averaging, conditional structures, or nonstandard decompositions are required.

3.9. Final Note on Geometry

The argument above selects an exponential tail because the multiplicative structure of sieve deviations becomes additive under logarithmic scaling, naturally leading to exponential behavior at the level of aggregated remainders.

This choice is therefore canonical within the present framework. Nevertheless, there is no *a priori* reason that the underlying geometry must be strictly exponential in all settings. Alternative functional families, such as hyperbolic profiles (e.g. \sinh^2 , or, in more structured contexts, Jacobi elliptic functions described by Abramowitz and Stegun [8] or Lambert-W transforms described by Corless et al. [9])—may arise when the interaction between local constraints and global renormalization exhibits cyclic or quasi-periodic behavior across scales. A problem as rigid as Goldbach’s conjecture does not probe sufficiently many degrees of freedom to distinguish most such higher-order geometric features; however, a derivation that explicitly incorporates Euler-cap effects may provide the additional structure needed to do so.

Accordingly, even a complete proof of the Goldbach specialization would not, by itself, uniquely determine the full geometric form of the tail factor T beyond the leading order. More refined distinctions are expected to emerge only in problems involving additional structural parameters or higher-dimensional configurations.

Finally, standard geometric operations, such as angle addition and subtraction, may be formulated abstractly in terms of set operations on the underlying candidate classes C_p , providing a natural algebraic interpretation of curvature interactions within the framework.

Remark. PCGH is motivated by structural considerations and by the behaviour of explicit sieve computations, but is not offered as a theorem. Its purpose is to supply a coherent organizing principle for the remainder phenomena across many prime-like configurations. In the next section, the framework is specialized to the Goldbach/HL–Windowed setting and a concrete conjecture tailored to that problem is derived. Eventual proof of such an individual conjecture should not be interpreted as a universal confirmation of the hypothesis itself.

4. Specialization to the Goldbach / HL–Windowed Geometry

This section instantiates PCGH 1 in the classical Goldbach setting. The aim is to obtain an explicit analytic form of the remainder functional

$$\varepsilon(2n; L) := G(2n; M) - \mathring{G}^{HL}(2n; M),$$

where $G(2n; M)$ denotes the measured number of Goldbach pairs in the window $|m| \leq M$, $\mathring{G}^{HL}(2n; M)$ is the HL–A prediction restricted to the same window, and L is a scale parameter governing the geometry of the remainder.

At this stage L is treated as an abstract scale; its concrete identification in terms of M will be justified below. As in Section 3, all remainder behavior is organized through the multiplicative decomposition

$$R^{HL}(2n; L) = H^{HL}(2n; L)H_2^{HL}(2n; L)T^{HL}(2n; L)T_2^{HL}(2n; L).$$

4.1. Hardy–Littlewood Conjecture A for Goldbach Windows

In order to apply PCGH to the Goldbach problem, we must first select an appropriate predictor function.

Hardy and Littlewood proposed the following asymptotic predictor Hardy–Littlewood Conjecture A (HL–A) for Goldbach representations:

$$2C_2\mathfrak{S}(2n)\frac{2n}{\log n}.$$

In previous work [5], this predictor was adapted to long intervals by replacing the global density factor with a general weighted sum,

$$2C_2\mathfrak{S}(2n) \sum_{m \in I^{par}(2n; M)} \omega(n-m)\omega(n+m).$$

Motivated by this formulation, let the *Hardy–Littlewood Windowed* predictor (HL–Windowed) restrict the same structure to a local Goldbach window,

$$\hat{G}^{HL}(2n; M) := 2C_2 \mathfrak{S}(2n) \sum_{m \in I^{par}(2n; M)} \omega(n-m)\omega(n+m),$$

as specified in Definition 10. This represents a finite–window specialization of Hardy–Littlewood Conjecture A rather than a distinct conjecture.

Equivalently, for Goldbach configurations $(n-m, n+m)$ with local window $|m| \leq M$, the Hardy–Littlewood A predictor (including the standard circle–method correction) for *ordered* pairs may be written schematically as

$$2 \left(\prod_{p \in \mathbb{P} \setminus \{2\}} \left(1 - \frac{1}{(p-1)^2} \right) \right) \left(\prod_{\substack{p \in \mathbb{P} \setminus \{2\} \\ p \mid n}} \frac{p-1}{p-2} \right) \sum_{m \in I^{par}(n, M)} \omega(n-m)\omega(n+m),$$

where $\mathbb{P} \setminus \{2\}$ denotes the set of odd primes.

Here:

- the first product is the classical singular series, converging to the twin Hardy–Littlewood constant C_2 ;
- the second product represents the small– and medium–prime sieve residual arising from the obstruction $p \mid n$;
- the sum over m encodes the local density within the Goldbach window.

This HL–A predictor is a special case of the general predictive framework introduced previously, with irrelevant primes omitted once their associated local moduli exceed the window scale.

Remark. A naïve replacement of the Hardy–Littlewood weight $1/\log x$ by $li(x)/x$ is inappropriate in a windowed setting. The correct analogue requires integration over the window itself, which induces a continuous correction that must be paired with a corresponding window-normalized constant. When both effects are included, the resulting corrections largely cancel, explaining the empirical effectiveness of the classical Hardy–Littlewood normalization in finite windows.

Remark (Predictor Improvements). The framework developed here naturally suggests more accurate finite- n predictors, obtained by incorporating additional cutoff-dependent correction factors. Such refinements may improve numerical agreement over limited ranges.

The objective of the present work, however, is not to optimize finite-range accuracy, but to formulate a conjecture grounded in a predictive structure that has withstood more than a century of scrutiny. For this reason, the classical Hardy–Littlewood form, deferring any modified predictors, whose finite- n behaviour remains comparatively untested, to potential future studies.

Hardy–Littlewood Conjecture A proposes an asymptotic predictor for Goldbach representations, but no proof is known that this predictor converges to the measured Goldbach count G . Empirical investigations, including the computational study in [5], nevertheless indicate strong agreement over large ranges.

Remark (Euler Cap In the HL–A Normalization).

The Hardy–Littlewood prediction

$$\hat{G}^{HL}(2n) = 2C_2 \mathfrak{S}(2n) n \omega^2(n)$$

necessarily entails the presence of a saturation mechanism: the total Goldbach mass at level $2n$ cannot grow without bound as the size of a local window increases. In this sense, HL–A implicitly assumes the existence of a cap. In the present framework, this saturation is known explicitly via what we call the *Euler cap*, which bounds the total mass by a fixed multiple of $2n\omega^2(n)$.

Accordingly, while a local proxy of the form

$$\sum_{m \in I^{par}(n, M)} \omega(n-m)\omega(n+m) \approx 2M\omega^2(n)$$

is accurate whenever $M = o(n)$, it must necessarily saturate as M approaches n .

The Euler cap does not affect the geometric structure developed in this paper. Its role becomes relevant only when analytic proxy formulas are invoked, where it suppresses spurious growth, such as additional logarithmic factors, that would otherwise arise for large $\alpha = M/n$.

A detailed discussion of the Euler cap and its role in the HL–A normalization is given in previous work [5]. Throughout the present paper, the analysis remains in the pre-cap regime, where M is sufficiently small that the linear local proxy remains valid.

Remark (Goldbach Notation and Conventions).

Throughout the Goldbach section:

- $G(2n; M)$ denotes the number of Goldbach pairs $(n-m, n+m)$ with $m \in I^{par}(n, M)$;
- $\hat{G}^{HL}(2n; M)$ denotes the Hardy–Littlewood A circle–method prediction restricted to the same window, dropping unnecessary primes;

- $\mathcal{Q}(2n; L)$ denotes the effective local modulus;
- $\mathfrak{S}^C(2n; L)$ is the complement truncated semiprime correction factor;
- $\hat{\Omega}_{\text{prime}}(2n; L)$ is the renormalized prime curvature factor defined in (114).

Unless otherwise stated, Goldbach pairs are counted as ordered pairs $(n - m, n + m)$; for unordered pairs the factor 2 is omitted.

4.2. The Goldbach Configuration and Natural Scale

For fixed even $2n$, the configuration set is

$$C^{(i)} = \{(n - m, n + m) : m \in I^{\text{par}}(n, M)\}.$$

This is a symmetric block of size $2M$; its natural scale is

$$L := \sqrt{2M}.$$

This scale is not arbitrary: it is the unique one for which (i) the classical HL–A prediction stabilizes, (ii) residue-class distortions from small primes remain controlled, and (iii) extremal deviations grow no faster than $O(L \log n / \log \log n)$. The scale $L := \sqrt{2M}$ therefore plays the same role here that the block scale $\Lambda(C^{(i)})$ played in the general hypothesis.

4.3. Prime–Like Sets and the Scale Function $\beta(p)$

In the Goldbach setting, all arithmetic structure is generated by the genuine primes. Accordingly, a single prime–like set

$$S = P,$$

with focus on the construction of an appropriate scale function

$$\beta(p), \text{ with } p \leq P_0(L_i) \iff \beta(p) \leq L_i$$

is compatible with the geometry of a symmetric Goldbach window.

For a fixed even integer $2n$, the local configuration is considered

$$m \in I^{\text{par}}(2n; M),$$

to seek to understand which primes exert structurally significant influence on the distribution of Goldbach pairs within this window.

Empirical inspection and heuristic analysis indicate that the relevant residue structure is governed by a family of rapidly growing *effective moduli*. The lower cutoff prime is first introduced

$$p_{\min}(n) := \begin{cases} 3, & 3 \nmid n, \\ 5, & 3 \mid n. \end{cases}$$

This definition reflects a peculiarity of the Goldbach configuration: when $3 \nmid n$, exactly one of $n - m$ or $n + m$ is divisible by 3 for every odd m . In that case, the prime 3 acts as a fixed sieving prime across the entire window and does not contribute to local variation. Only when $3 \mid n$ does the prime 3 participate nontrivially in the local residue structure, necessitating its inclusion among the small primes. This behavior is unique to Goldbach-type configurations and motivates the definition of $p_{\min}(n)$.

For each odd prime p , recall from Definition 5 that the effective local modulus attached to the even integer $2n$ is

$$\mathcal{Q}_p(n) := \prod_{\substack{q \in \mathbb{P}_{\text{eff}}(n) \\ q \leq p}} (q - 1),$$

where $p_{\min}(n) \in \{3, 5\}$ is chosen according to the condition $3 \mid n$. Thus $\mathcal{Q}_p(n)$ measures the size of the residue-class orbit induced by all odd primes $q < p$, with each factor $(q - 1)$ reflecting the effective sieving density at modulus q .

In contrast to the raw primorial $\prod_{q < p} q$ suggested by a naive application of the Chinese Remainder Theorem, the use of $(q - 1)$ aligns with the Hardy–Littlewood singular-series geometry: for Goldbach-type problems, each prime contributes according to its admissible residue classes rather than its full modulus.

The rapid growth of $\mathcal{Q}_p(n)$ induces a natural and empirically accurate cutoff criterion for the influence of a prime p on a Goldbach window. In writing

$$\mathcal{Q}_p(n) \leq L \quad \iff \quad \mathcal{Q}_p(n) \leq L := \sqrt{2M},$$

it is shown that primes satisfying this condition exert coherent local influence on the window $m \in I^{par}(n, M)$. Primes beyond this threshold affect the count only through smooth, aggregate fluctuations and are therefore absorbed into the medium-prime tail $T^{HL}(2n; L)$ in the remainder decomposition.

This construction yields a Goldbach-specific realization of the general scale separation framework developed in Section 3. The effective moduli $\mathcal{Q}_p(n)$ determine the appropriate small-prime cutoff $P_0(L)$, and hence the scale function $\beta(p)$, in a manner consistent with both the Hardy–Littlewood prediction and the observed geometry of Goldbach deviations.

4.4. The Small-Prime Factors $H^{HL}(2n; L)$ and $H_2^{HL}(2n; L)$

Only a few small primes appear in the effective modulus $\mathcal{Q}_p(n)$, but these primes impose the dominant, non-averaging residue-class constraints on the Goldbach window. It is therefore natural to expect the finite-scale factors H^{HL} and H_2^{HL} to mirror the respective roles of the prime-pair constant C_2 and the local correction factor \mathfrak{S} in the Hardy-Littlewood formulation.

The combined small-prime contribution $H^{HL}(2n; L)H_2^{HL}(2n; L)$ is given by the finite product

$$H^{HL}(2n; L)H_2^{HL}(2n; L) := \prod_{\substack{p \in \mathbb{P}_{\text{eff}}(n) \\ \mathcal{Q}_p(n) \leq L}} h(p; n)h_2(p; n),$$

where, in the Goldbach setting, the local factor satisfies

$$h(p; n)h_2(p; n) = \begin{cases} p-2, & \text{if } n \text{ lies in an admissible residue class modulo } p, \\ p-1, & \text{if } p \mid n. \end{cases}$$

This decomposition motivates separating the residue-averaged and configuration-dependent components. The *global small-prime factor* is defined as

$$H^{HL}(2n; L) := \prod_{\substack{p \in \mathbb{P}_{\text{eff}}(n) \\ \mathcal{Q}_p(n) \leq L}} (p-2),$$

and the *local correction factor* as

$$H_2^{HL}(2n; L) := \prod_{\substack{p \in \mathbb{P}_{\text{eff}}(n) \\ p \mid n \\ \mathcal{Q}_p(n) \leq L}} \frac{p-1}{p-2}.$$

The latter quantity is the finite-scale analogue of the classical Hardy-Littlewood semiprime correction.

Therefore the notation \mathfrak{S} is extended to the cutoff scale L by setting

$$\mathfrak{S}(2n; L) := \mathfrak{S}_{\text{head}}^{\mathcal{Q}}(2n; L) := H_2^{HL}(2n; L).$$

Likewise, $\mathfrak{S}_{\text{head}}^{\mathcal{C}}(2n; L)$ uses the condition $p \nmid n$, while $\mathfrak{S}_{\text{head}}^{\cdot}(2n; L)$ removes the divisibility restriction.

Both $H^{HL}(2n; L)$ and $\mathfrak{S}_{\text{head}}^{\mathcal{Q}}(2n; L)$ are finite Euler products. Indeed, the cutoff condition $\mathcal{Q}_p(n) \leq L$ is satisfied only by the first few primes, since $\mathcal{Q}_p(n)$ grows super-factorially with p . As a result, the cutoff index increases only when L grows by several orders of magnitude. In this context, rapid convergence is not meaningful: the contribution of the small-prime factors is determined entirely by very small primes.

Consequently, $H^{HL}(2n; L)$ and $\mathfrak{S}_{\text{head}}^{\mathcal{Q}}(2n; L)$ are piecewise constant functions of L , changing only at isolated “phase transition” points where a new prime p enters the small-prime range. Between these transitions,

the small-prime contribution is fixed, and the variation in the remainder geometry is driven almost entirely by the tail factor $T^{HL}(2n; L)$.

At this point it is worth noting the following relationship,

$$\begin{aligned}
H^{HL}(2n; L)H_2^{HL}(2n; L) &= H^{HL}(2n; L)\mathfrak{S}_{\text{head}}^{\mathcal{Q}}(2n; L). \\
&= \frac{H^{HL}(2n; L)\mathfrak{S}_{\text{head}}^{\mathcal{Q}}(2n; L)\mathfrak{S}_{\text{head}}^{\mathcal{Q}, \mathcal{C}}(2n; L)}{\mathfrak{S}^{\mathcal{Q}, \mathcal{C}}(2n; L)}, \\
&= \frac{H^{HL}(2n; L)\mathfrak{S}_{\text{head}}^{\mathcal{Q}, \mathcal{C}}(2n; L)}{\mathfrak{S}_{\text{head}}^{\mathcal{Q}, \mathcal{C}}(2n; L)}, \\
&= \left(\prod_{\substack{p \in \mathbb{P}_{\text{eff}}(n) \\ \mathcal{Q}_p(n) \leq L}} p - 2 \right) \left(\prod_{\substack{p \in \mathbb{P}_{\text{eff}}(n) \\ p \mid n \\ \mathcal{Q}_p(n) \leq L}} \frac{p-1}{p-2} \right) / \mathfrak{S}_{\text{head}}^{\mathcal{Q}, \mathcal{C}}(2n; L). \\
&= \left(\prod_{\substack{p \in \mathbb{P}_{\text{eff}}(n) \\ \mathcal{Q}_p(n) \leq L}} p - 1 \right) / \mathfrak{S}_{\text{head}}^{\mathcal{Q}, \mathcal{C}}(2n; L).
\end{aligned}$$

This leads to more natural notation for header terms,

$$H^{HL}(2n; L)H_2^{HL}(2n; L) = \frac{\mathcal{Q}(2n; L)}{\mathfrak{S}_{\text{head}}^{\mathcal{Q}, \mathcal{C}}(2n; L)}$$

4.5. The Medium-Prime Tail $T^{HL}(2n; L)$

For the Goldbach specialization, the medium primes are those with $\mathcal{Q}_p(n) \leq L$. Their contribution to the remainder is encoded in the tail factor

$$T^{HL}(2n; L) := \prod_{\substack{p \in \mathbb{P}_{\text{eff}}(n) \\ \mathcal{Q}_p(n) \leq L}} t(p; n)^{L/\mathcal{Q}_p(n)},$$

where $\mathcal{Q}_p(n)$ is the effective local modulus from Definition 5. In the HL-A Goldbach setting the PCGH tail multiplier is

$$t(p; n) = p - 2,$$

reflecting the loss of two residue classes modulo p for admissible Goldbach pairs.

In the abstract PCGH geometry the tail is re-expressed as an exponential curvature term $T(L) = (\Omega\Theta(L))^L$.

For the HL-A setting it is convenient to separate a single global constant and a finite renormalization

factor that depends on $2n$ and the cutoff L .

Fix the base prime $q_{\min} = 5$ and recall the auxiliary moduli $\mathcal{Q}_p^{(5)}$ from Definition 5. The *prime curvature constant* by the convergent Euler product is defined as

$$\Omega_{\text{prime}} := \prod_{\substack{p \in \mathbb{P} \\ p \geq 5}} (p-2)^{1/\mathcal{Q}_p^{(5)}}.$$

This constant depends only on the Goldbach prime–pair geometry and is independent of n and M .

For each even integer $2n$ and Euler–cap admissible window size M , the *renormalized curvature factor* is defined as

$$\hat{\Omega}_{\text{prime}}(2n; L) := \Omega_{\text{prime}}^{\kappa(n)} \prod_{\substack{p \in \mathbb{P}_{\text{eff}}(n) \\ \mathcal{Q}_p(n) \leq L}} (p-2)^{-1/\mathcal{Q}_p(n)},$$

where $\kappa(n) \in \{\frac{1}{2}, 1\}$ is a bounded exponent accounting for the choice $p_{\min}(n) \in \{3, 5\}$ in Definition 5.

When $3 \nmid n$ one has $p_{\min}(n) = 5$ and $\kappa(n) = 1$; when $3 \mid n$ the extra contribution from the prime 3 is absorbed into the exponent $\kappa(n) = \frac{1}{2}$. The precise value of $\kappa(n)$ is not important for the bounds that follow; only that it is uniformly bounded and depends at most on the residue class of $n \bmod 3$.

A straightforward rearrangement of the Euler products then shows that

$$T^{HL}(2n; L) = \left(\hat{\Omega}_{\text{prime}}(2n; L) \right)^L.$$

Thus all dependence on the medium primes above the cutoff is captured by a single scale–invariant curvature constant Ω_{prime} , together with a finite renormalization determined by the small primes below the cutoff and the residue of n modulo 3. Between successive renormalization scales (when a new prime enters the small–prime range) the map $L \mapsto \hat{\Omega}_{\text{prime}}(2n; L)$ is constant, and all variation in the tail is carried by the power L .

Remark (Role of the Renormalized Curvature Factor). The factor $\hat{\Omega}_{\text{prime}}(2n; L)$ is not introduced to create additional growth in the remainder term, but to ensure that the *small–prime block together with its tail* behaves smoothly as the cutoff moves.

Corollary A.1 shows that, when a new prime p_* crosses the cutoff, its factor $(p_* - 2)$ is gradually transferred from the tail into the small–prime product: just before the transition it appears in the tail with a fractional exponent $L/\mathcal{Q}_{p_*}(n) \in (0, 1)$, and at the transition it is fully absorbed into $H^{HL}(2n; L)$. The renormalization in $\hat{\Omega}_{\text{prime}}(2n; L)$ is chosen so that this transfer leaves the product

$$H^{HL}(2n; L) T^{HL}(2n; L)$$

continuous as a function of L , even though $H^{HL}(2n; L)$ and $T^{HL}(2n; L)$ separately have discrete jumps at the renormalization scales.

As more primes are shifted from the tail into the small-prime base, the finite renormalization in (114) cancels an increasing portion of the infinite product defining the curvature constant. Along any sequence of scales passing through successive cutoffs this drives

$$\hat{\Omega}_{\text{prime}}(2n; L) \longrightarrow 1,$$

while $\log T^{HL}(2n; L)$ retains exactly the leading $L/\log\log L$ behaviour from Lemma A.2, in the sense that

$$\frac{\log T^{HL}(2n; L)}{L/\log\log L} = 1 + o(1) \quad (L \rightarrow \infty),$$

for each fixed residue class of $n \bmod 3$. Thus $\hat{\Omega}_{\text{prime}}(2n; L)$ serves as a renormalizing bridge: it enforces continuity of the curvature geometry for the $H^{HL} \cdot T^{HL}$ block across phase transitions, while its deviation from 1 remains uniformly bounded and asymptotically negligible at the scale of the overall remainder.

Remark (Alternate Notation). While the usage of $\hat{\Omega}_{\text{prime}}(2n; L)^L$ allows for analytic analysis, the exponent really just reminds one it is a curvature. A more abbreviated notation:

$$\Xi(2n; L) := \hat{\Omega}_{\text{prime}}(2n; L)^L.$$

eliminates the need to record the exponent separately.

4.6. The Residual Factor $T_2^{HL}(2n; L)$

Conceptually, the factor $T_2(2n; L)$ collects all residue-dependent corrections coming from primes p with

$$\mathcal{Q}_p(n) > L := \sqrt{2M},$$

that is, primes whose effective moduli are so large that they do not create coherent structure inside the Goldbach window $m \in I^{par}(n, M)$, but still distinguish cases such as $p \mid n$ versus $p \nmid n$.

Without loss of generality, and to make the dependence on $3 \mid n$ transparent, consider first the case $3 \mid n$, so that $p_{\min}(n) = 3$ and

$$\mathcal{Q}_p(n) = \mathcal{Q}_p^{(3)} := \prod_{\substack{q \in \mathbb{P} \\ 3 \leq q < p}} (q-1),$$

with the complementary case $3 \nmid n$ handled with slight adjustment, the same applies by replacing $\mathcal{Q}_p^{(3)}$ with $\mathcal{Q}_p^{(5)}$.

At the level of a single prime, there is no reason to expect $T_2(2n; L)$ to be numerically close to 1: for the first "very small" prime p_* above the cutoff, the residue correction at that prime multiplies the tail by

$$\frac{p_* - 1}{p_* - 2} \quad \text{when } p_* \mid n.$$

For instance, if $p_* = 11$ then this local factor is $10/9 \approx 1.11$. Globally, however, the corresponding change in $\log T_2$ is attenuated by the large modulus $\mathcal{Q}_{p_*}(n)$, so the net effect on the tail at scale L is already small.

The key point is that this local correction is already included in the definition of the effective moduli. A naive primorial geometry would use

$$\mathcal{Q}_p^{prim} := \prod_{\substack{q \in \mathbb{P}_{\text{eff}}(n) \\ q < p}} q,$$

so that the loss of one residue class at p_* would appear in the tail as an extra factor of order $(p_* - 1)/(p_* - 2)$. Instead, Definition 5 replaces each q by $(q - 1)$, so that in the $3 \mid n$ case

$$\mathcal{Q}_p^{(3)} = \prod_{\substack{q \in \mathbb{P} \\ 3 \leq q < p}} (q - 1).$$

This $(q - 1)$ -based scaling makes the main tail $T^{HL}(2n; L)$ slightly larger than in the primorial model, by essentially the same amount that $T_2(2n; L)$ contributes.

Moreover, $T_2(2n; L)$ only contains residue adjustments for primes with $\mathcal{Q}_p(n) > L$, and these adjustments always *reduces* the effective tail when $p \mid n$. Since the choice of $\mathcal{Q}_p(n)$ makes $T^{HL}(2n; L)$ conservative even in the semiprime-saturated case where every small prime divides n , the overall remainder bound

$$|\varepsilon(2n, L)| \leq 2 \frac{\mathcal{Q}(2n; L)}{\mathfrak{S}_{\text{head}}^{\mathcal{Q}, C}(2n; L)} \left(\tilde{\Omega}(2n; L) \right)^L$$

remains valid if one simply sets

$$T_2^{HL}(2n; L) \equiv 1$$

in the geometry. The corresponding statements for the case $3 \nmid n$ follow by the same argument with $\mathcal{Q}_p^{(5)}$ in place of $\mathcal{Q}_p^{(3)}$.

In summary, T_2 is not numerically 1 at the level of a single prime, but its effect is (i) diluted by the large moduli $\mathcal{Q}_p(n)$ and (ii) pre-compensated by the $(q - 1)$ -based definition of $\mathcal{Q}_p(n)$. For the purpose of conservative upper bounds on $|\varepsilon(2n, L)|$, it is therefore safe and conceptually cleaner to set $T_2^{HL}(2n; L) = 1$.

4.7. Specialized Remainder Decomposition

Collecting the small-prime factors $H^{HL}(2n; L)$, the semi-local correction $\mathfrak{S}_{\text{head}}^{\mathcal{Q}, \mathcal{C}}(2n; L)$, and the medium-prime tail $T^{HL}(2n; L)$ from the preceding subsections, and taking $T_2^{HL}(2n; L) \equiv 1$ as justified above, the Goldbach remainder envelope takes the form

$$\begin{aligned} R^{HL}(2n; L) &:= 2 \frac{\mathcal{Q}(2n; L)}{\mathfrak{S}_{\text{head}}^{\mathcal{Q}, \mathcal{C}}(2n; L)} \left(\hat{\Omega}_{\text{prime}}(2n; L) \right)^L \quad L := \sqrt{2M}, \\ &= 2 \frac{\mathcal{Q}(2n; L) \Xi(2n; L)}{\mathfrak{S}_{\text{head}}^{\mathcal{Q}, \mathcal{C}}(2n; L)} \quad \Xi(2n; L) := \left(\hat{\Omega}_{\text{prime}}(2n; L) \right)^L. \end{aligned}$$

where $\hat{\Omega}_{\text{prime}}(2n; L)$ is the renormalized curvature factor defined in (114), built from the universal constant Ω_{prime} of (113). The factor 2 reflects the standard convention of counting *ordered* Goldbach pairs $(n - m, n + m)$; for unordered pairs this factor is omitted.

The expression (127) will serve as the exact geometric remainder envelope in the Goldbach specialization of the Prime Curvature framework. Its conjectural role is stated explicitly in the next section.

Remark (Primorial vs. $(q - 1)$ scaling). It is instructive to contrast the $(q - 1)$ -based moduli $\mathcal{Q}_p(n)$ with the more naive "primorial" choice $\prod_{3 < q < p} q$. In a primorial model one would retain an explicit residual tail factor T_2 to account for residue corrections at primes just above the cutoff. The resulting bound can be tuned to match extremal deviations, but it is less well behaved away from those spikes.

By working instead with the $(q - 1)$ -scaled moduli $\mathcal{Q}_p(n)$ and absorbing the corresponding cost into $T^{HL}(2n; L)$, the residual factor T_2 becomes superfluous: its effect is already accounted for in the conservative choice of tail exponent, and the simplified envelope $R^{HL}(2n; L)$ still captures the correct $L/\log\log L$ scale from Lemma A.2. In this sense the $(q - 1)$ scaling trades a small loss of extremal sharpness for a gain in *uniform coverage*, which is precisely what is needed for a remainder geometry intended to control $\varepsilon(2n; L)$ pointwise rather than only at its extremes.

4.8. Prime Curvature Geometry Conjecture For Goldbach

We now state the specialization of the general curvature framework to the Goldbach configuration. The full conjectures are expressed using the exact remainder functional $R^{HL}(2n; L)$ derived in this section; bounding envelope forms, when used, serve only as analytic tools and are not part of the statement.

Conjecture 1 (Prime Curvature Geometry Conjecture for Goldbach (PCGC–Goldbach)).

For every even integer $2n \geq 4$ and every admissible window size M within the Euler cap, the Goldbach pair counts satisfy

$$\left| G(2n; M) - \mathring{G}^{HL}(2n; M) \right| \leq R^{HL}(2n; L), L := \sqrt{2M},$$

where the remainder envelope $R^{HL}(2n; L)$ is defined in (127).

Remark (Strength Relative to HL–Windowed). As discussed in the entrance criteria for PCGH, the curvature framework developed in this paper is formulated under a pointwise asymptotic assumption on the main term. In the Goldbach setting, this takes the form

$$G(2n; M) = \mathring{G}^{HL}(2n; M)(1 + o(1)),$$

which is obtained via control of the normalized remainder $R^{HL}(2n; M)/\mathring{G}^{HL}(2n; M)$.

This assertion is strictly stronger than the classical Hardy–Littlewood A conjecture, which governs averaged behaviour and does not claim pointwise convergence of individual Goldbach counts. Accordingly, it is logically possible for Hardy–Littlewood A to hold while PCGC–Goldbach fails.

The role of PCGH is therefore not to justify the Hardy–Littlewood main term, but to isolate the additional geometric regularity required for pointwise control of deviations once such a main term is assumed.

Remark (On Power–Envelope Variants). One might consider a family of envelope inequalities of the form

$$\left| G(2n; M)^A - \mathring{G}^{HL}(2n; M)^A \right| \leq R^{HL}(2n; L)^A, L := \sqrt{2M}, A \geq 1.$$

For $A \neq 1$, this implicitly imposes additional shape constraints on the deviation $G - \mathring{G}^{HL}$ (e.g. penalizing large excursions differently) and, absent an independent structural justification, effectively introduces a tunable parameter into the conjecture. Accordingly, PCGC–Goldbach is stated with the canonical choice $A = 1$. Nevertheless, it is informative to test (131) empirically for other values of A , since such measurements may suggest a sharper (falsifiable) refinement if a natural value of A emerges from the data.

Remark (From Pair Counts to Densities). The quantities $G(2n; M)$ and $\mathring{G}^{HL}(2n; M)$ are best interpreted not as fundamental combinatorial objects, but as finite–window measurements of an underlying Goldbach pair density. From this perspective, the role of HL–Windowed is not to assert the existence of individual pairs, but to provide an asymptotic model for this density.

The conjectures formulated in this paper therefore concern the boundedness and convergence of these density measurements, rather than pointwise pair counts. This places the Goldbach problem in direct

analogy with the theory of singular series, where convergence and normalization, rather than exact enumeration, are the primary analytic objects.

Remark (Short-Interval Obstructions). Even with the use of a windowed semilocal correction term $\mathfrak{S}(2n; M)$, the Prime Curvature Geometry framework does not preclude the existence of isolated short intervals in which all admissible Goldbach pairs are cancelled by divisibility constraints arising from medium or large primes that do not divide $2n$. Such primes do not contribute to $\mathfrak{S}(2n; M)$, and therefore their effect is invisible to semilocal corrections tied solely to the target value.

These obstructions do not represent a failure of the conjecture. Any cancellation pattern constructed entirely from medium and large primes can be extended only over a finite range, on the order of the number of terms of a primorial, before exhausting the supply of primes capable of aligning obstructively without contributing to the correction structure. Consequently, such phenomena cannot persist once the interval reaches a scale at which asymptotic windowed weights accurately reflect prime density, at which point cancellation effects necessarily re-enter the analytic correction terms.

This observation underscores that PCGC--Goldbach controls averaged behaviour and sufficiently large windows, but should not be interpreted as predicting a minimal short--interval scale for convergence. In particular, intervals of size $M = n\alpha(n)$ with $\alpha(n) = O\left(\frac{\log^k n}{n}\right)$ may exhibit out of bounds predictions unless appropriate window--dependent weighting is employed.

4.9. Bounding Envelope (Derived)

The Goldbach deviation satisfies the abstract bound

$$|\varepsilon(2n; L)| \leq R^{HL}(2n; L),$$

where $R^{HL}(2n; L)$ is the exact geometric remainder envelope defined in (127).

For analytic comparison and for identifying admissible growth regimes, it is convenient to work with an explicit bounding envelope derived from Lemma A.2. With $L = \mathcal{Q}(2n; L)$, define

$$\hat{R}^{HL}(2n; L) := \frac{c(2n; L)L}{\mathfrak{S}_{\text{head}}^{\mathcal{Q}, \mathcal{C}}(2n; L)},$$

where the correction factor $c(2n; L)$ is given by the explicit product formula (A.25).

By Lemma A.4, the factor $c(2n; L)$ satisfies

$$c(2n; L) \geq 2, c(2n; L) \searrow 2 \quad (L \rightarrow \infty),$$

uniformly in n . In particular, $c(2n; L)$ decreases monotonically as the effective modulus cutoff increases, reflecting the progressive absorption of medium-prime curvature into the exact envelope.

Combining this with the definition of R^{HL} in (127), Lemma A.2 implies that the bounding envelope $\hat{R}^{HL}(2n; L)$ is conservative and asymptotically sharp:

$$R^{HL}(2n; L) \leq \hat{R}^{HL}(2n; L) = R^{HL}(2n; L) \left(1 + o_L(1)\right),$$

uniformly over the admissible parameter ranges.

Remark (Logical Status of Derived Bounds). Any theorem or lemma derived conditionally from PCGC–Goldbach may be viewed as a weaker conjecture obtained by relaxing the exact geometric envelope R^{HL} to a dominating bound. Such weakened statements retain substantial analytic content; sufficient for growth estimates, admissible window analysis, and comparison arguments; while discarding detailed geometric assumptions that are not essential for these purposes.

The bounding envelope \hat{R}^{HL} is *not* part of the conjectural framework itself. It is a derived analytic tool used for comparison arguments, growth estimates, and the identification of admissible window scales. No claim is made here regarding how frequently the true deviation approaches saturation of this bound.

In summary, PCGC–Goldbach isolates a single geometric envelope governing the deviation of windowed Goldbach counts from the Hardy–Littlewood prediction. The remainder is controlled by:

- explicit small-prime structure through the effective moduli $\mathcal{Q}_p(n)$;
- finite divisibility corrections through $\mathfrak{S}_{\text{head}}^{\mathcal{Q}, \mathcal{C}}(2n; L)$; and
- a renormalized exponential curvature factor associated with the medium-prime tail.

The conjecture itself is stated entirely in terms of the exact envelope R^{HL} ; the bounding envelope \hat{R}^{HL} serves solely as an analytic surrogate.

5. Conclusion

This paper develops a finite-scale geometric framework for analyzing remainder terms in additive prime problems and formulates a precise curvature bound for Goldbach representations. The central structural feature is the separation of arithmetic effects into three distinct components: a finite small-prime factor governed by effective moduli, a divisibility correction capturing genuine discontinuities in the central parameter n , and a renormalized medium-prime tail that controls scale-dependent fluctuations.

In the Goldbach setting, the effective moduli $\mathcal{Q}_p(n)$ induce a sequence of scale thresholds at which new primes enter the small-prime range. Each such transition produces a discrete change in the finite Euler product, while the associated tail factor interpolates smoothly across the cutoff. This mechanism explains how piecewise-constant local structure can coexist with globally stable remainder envelopes without recourse to global averaging or asymptotic smoothing.

PCGC–Goldbach asserts that the deviation from the Hardy–Littlewood prediction is uniformly bounded by an explicit geometric envelope constructed from these components. The conjecture is stated in its full, non-proxy form; simplified bounding envelope expressions are introduced to simplify calculations and to elucidate growth rates and threshold behaviour and play no role in the formulation of the bound itself.

A key feature of the framework is that divisibility–driven discontinuities are retained explicitly, rather than being absorbed into error terms. This preserves honesty at finite scales and avoids spurious smoothness in regimes where arithmetic structure is dominant. At the same time, scale–induced discontinuities are neutralized by the renormalized tail, yielding a stable geometric description across all admissible window sizes.

The curvature formulation presented here provides a clean structural alternative to classical sieve–based remainder analysis. It isolates the precise sources of non–averaging behaviour and organizes them into a scale–consistent bound. Subsequent work will address numerical calibration of the curvature envelope and extensions to related additive problems, including mixed prime–semiprime configurations, where multiple prime–like sets interact at finite scales.

5.1. Consequences and Outlook

The curvature framework developed here is intentionally local: it separates the Goldbach geometry into a small-prime block, a finite divisibility correction, and a renormalized medium-prime tail, with an explicit remainder functional $R^{HL}(2n; L)$. PCGC–Goldbach asserts that this remainder envelope is globally valid for all admissible windows M and all $2n \geq 4$.

In subsequent work on reductions, this geometric envelope will be used in a attempt to translate PCGC–Goldbach into more classical consequences. Under the conjecture one recovers, for example:

- asymptotic agreement between measured Goldbach counts and the Hardy–Littlewood prediction across a wide range of logarithmic window scales where asymptotic weighting applies;

- existence of prime pairs in intervals of length $M = n\alpha(n)$ throughout regimes in which windowed asymptotic weights accurately reflect prime density, including scales where $\alpha(n)$ approaches polylogarithmic order;
- and a pathway toward strengthening short–interval Goldbach results by improving the curvature threshold at which such windowed asymptotic control becomes effective.

If successful, these reductions will show that PCGC–Goldbach is not merely a reformulation of circle–method heuristics, but a genuinely geometric hypothesis whose validity propagates through a broad class of Goldbach–type statements once the appropriate analytic regime is reached.

Beyond Goldbach, the same renormalization mechanism applies to other additive problems whose remainder terms are governed by Hardy–Littlewood singular series. PCGH indicates that, once a curvature envelope is established for a given configuration, one may systematically transfer it into asymptotic predictions and short–interval existence results for the corresponding prime patterns. Developing these extensions, and sharpening the curvature thresholds in the Goldbach case, is delegated for future papers.

5.2. Conceptual Motivation

The geometric organization adopted in this paper was shaped in part by structural ideas encountered outside analytic number theory. In particular, several mechanisms that arise naturally in energy–based and hierarchical learning frameworks provided useful conceptual guidance.

In energy–based models [10] by LeCun et al., the introduction of new degrees of freedom reshapes the effective energy landscape, requiring a global rebalancing of previously stable configurations. Analogously, in the present framework, the entry of a new prime into the small–prime range induces a discrete renormalization step, while the associated tail factor compensates to preserve global stability.

Related phenomena appear in large–scale optimization methods [11][12] by Bottou et al., where local updates propagate globally only after crossing scale–dependent thresholds. Hierarchical representations in document and graph processing [13] by Haffner and LeCun exhibit similar behavior: discrete structural augmentations produce phase–like transitions in global representations.

These parallels are not intended as formal analogies or sources of mathematical justification. Rather, they reflect a shared structural pattern, namely, that global behaviour can remain stable despite discrete,

scale-triggered updates to the underlying structure. Recognizing this pattern informed the separation of small-prime factors, divisibility effects, and curvature tails employed in this paper.

Appendix A. Additional Lemmas

A.1. Cutoff Interpolation via Geometric Tail

Lemma A.1 (Cutoff Interpolation Via Geometric Tail).

Let $\beta(p)$ be a strictly increasing scale function on the primes. Let $h(p) > 0$ and $t(p) > 0$ be local multiplicative factors.

Define the cutoff index $P_0(L)$ by

$$p \leq P_0(L) \iff \beta(p) \leq L,$$

and define

$$H(L) := \prod_{p \leq P_0(L)} h(p), \quad T(L) := \prod_{\beta(p) > L} t(p)^{L/\beta(p)}.$$

Let p_1 be the smallest prime with $\beta(p_1) > L$. Then as L increases through the cutoff value $L = \beta(p_1)$, the product $H(L)T(L)$ undergoes a multiplicative jump given by

$$\frac{H(\beta(p_1))T(\beta(p_1))}{H(\beta(p_1)^-)T(\beta(p_1)^-)} = \frac{h(p_1)}{t(p_1)}.$$

Proof. For $L < \beta(p_1)$, the prime p_1 belongs to the tail, so

$$T(L) = t(p_1)^{L/\beta(p_1)} \prod_{p > p_1} t(p)^{L/\beta(p)}.$$

Since $0 < L/\beta(p_1) < 1$, this contribution varies continuously in L .

At the cutoff $L = \beta(p_1)$, the exponent of $t(p_1)$ reaches 1, and the prime p_1 moves from the tail into the header. Thus

$$T(\beta(p_1)) = \prod_{p > p_1} t(p)^{\beta(p_1)/\beta(p)}, \quad H(\beta(p_1)) = h(p_1) \prod_{p < p_1} h(p).$$

Immediately before the cutoff,

$$H(\beta(p_1)^-) = \prod_{p < p_1} h(p), \quad T(\beta(p_1)^-) = t(p_1) \prod_{p > p_1} t(p)^{\beta(p_1)/\beta(p)}.$$

Taking the ratio gives

$$\frac{H(\beta(p_1))T(\beta(p_1))}{H(\beta(p_1)^-)T(\beta(p_1)^-)} = \frac{h(p_1)}{t(p_1)},$$

as claimed. \square

Corollary A.1 (Goldbach Cutoff Interpolation).

In the Goldbach setting, let the scale function be

$$\beta(p) := \mathcal{Q}_p(n),$$

and take the local factors

$$h(p) = t(p) := p - 2.$$

Then the geometric tail is given by

$$T^{HL}(2n; L) := \prod_{\substack{p \in \mathbb{P}_{\text{eff}}(n) \\ \mathcal{Q}_p(n) > L}} (p-2)^{L/\mathcal{Q}_p(n)}.$$

Across each cutoff transition $L = \mathcal{Q}_{p_1}(n)$, the factor $(p_1 - 2)$ passes from the tail into the small-prime product exactly as its tail exponent reaches 1. Since $h(p) = t(p)$, the cutoff interpolation introduces no discontinuity: the combined product $H^{HL}(2n; L)T^{HL}(2n; L)$ is continuous across all prime cutoffs.

A.2. Overall Bounding Envelope

Lemma A.2 (Overall Bounding Envelope).

Let $n \in \mathbb{N}$ and let $M \geq 1$ be Euler-cap admissible, and set $L := \sqrt{2M}$. Define the Goldbach remainder bounding envelope

$$\hat{R}^{HL}(2n; L) := \frac{c(2n; L)L}{\mathfrak{S}_{\text{head}}^{\mathcal{Q}, \mathcal{C}}(2n; L)},$$

where $c(2n; L)$ is defined by (A.25). Then $\hat{R}^{HL}(2n; L) \geq R^{HL}(2n; L)$ for all admissible L . Moreover, \hat{R}^{HL} cannot serve as an asymptotic proxy for R^{HL} : there is no function $\eta(2n; L) \rightarrow 0$ as $L \rightarrow \infty$ (along Euler-cap admissible scales) such that

$$R^{HL}(2n; L) = \hat{R}^{HL}(2n; L)(1 + \eta(2n; L))$$

for all sufficiently large admissible L . Equivalently, $R^{HL}(2n; L)/\hat{R}^{HL}(2n; L) \rightarrow 1$.

Proof. Fix an admissible L , and define the cutoff endpoints of its cell by

$$L_0 := \max \{\mathcal{Q}_p(n) : \mathcal{Q}_p(n) \leq L\}, L_1 := \min \{\mathcal{Q}_p(n) : \mathcal{Q}_p(n) > L\},$$

and set $c_i := c(2n; L_i)$ for $i \in \{0, 1\}$. Then $L \in (L_0, L_1]$, and on this cutoff cell the quantities L_0, L_1, c_0, c_1 are fixed.

By definition of \hat{R}^{HL} and the cutoff identities, then

$$R^{HL}(2n; L_i) = \hat{R}^{HL}(2n; L_i) = \frac{c_i L_i}{\mathfrak{S}_{\text{head}}^{\mathcal{Q}, \mathcal{C}}(2n; L_i)}, i \in \{0, 1\}.$$

By Corollary A.1, the envelope admits a continuous interpolation between successive cutoffs, and normalizing at L_1 we may write

$$R^{HL}(2n; L) = \frac{c_0 L_0}{\mathfrak{S}_{\text{head}}^{\mathcal{Q}, \mathcal{C}}(2n; L)} \left(\frac{c_1 L_1}{c_0 L_0} \right)^{L/L_1}.$$

Dividing by $\hat{R}^{HL}(2n; L) = \frac{c_0 L}{\mathfrak{S}_{\text{head}}^{\mathcal{Q}, \mathcal{C}}(2n; L)}$ gives

$$\rho(2n; L) := \frac{R^{HL}(2n; L)}{\hat{R}^{HL}(2n; L)} = \frac{L_0}{L} \left(\frac{c_1 L_1}{c_0 L_0} \right)^{L/L_1} = \gamma^{-1} e^{(\gamma-1)\omega},$$

where

$$\gamma := \frac{L}{L_1} \in \left(\frac{L_0}{L_1}, 1 \right], \omega := \log \left(\frac{c_1 L_1}{c_0 L_0} \right) > 0.$$

Differentiation yields

$$\frac{d\rho}{d\gamma} = \gamma^{-2} e^{(\gamma-1)\omega} (-1 + \gamma\omega),$$

so ρ has a unique stationary point at $\gamma_* := 1/\omega$, which is a minimum when $\gamma_* \in (\frac{L_0}{L_1}, 1)$. At this point,

$$\rho_{\min} = \rho(\gamma_*) = \omega e^{1-\omega}.$$

In particular, along any sequence of cutoff cells for which $\omega \rightarrow \infty$, we have $\rho_{\min} \rightarrow 0$, hence there exists an admissible sequence $L_* \rightarrow \infty$ (with L_* chosen at the within-cell minimizer) such that

$$\frac{R^{HL}(2n; L_*)}{\hat{R}^{HL}(2n; L_*)} = \rho(2n; L_*) \rightarrow 0.$$

On the other hand, at the endpoint $L = L_1$ we have $\rho(2n; L_1) = 1$.

Now suppose for contradiction that there exists $\eta(2n; L) \rightarrow 0$ such that $R^{HL}(2n; L) = \hat{R}^{HL}(2n; L)(1 + \eta(2n; L))$ for all sufficiently large admissible L . Then

$$\rho(2n; L) = \frac{R^{HL}(2n; L)}{\hat{R}^{HL}(2n; L)} = 1 + \eta(2n; L) \longrightarrow 1,$$

as $L \rightarrow \infty$, contradicting the admissible sequence $L_* \rightarrow \infty$ with $\rho(2n; L_*) \rightarrow 0$. Therefore no such η exists, and \hat{R}^{HL} is not an asymptotic proxy for R^{HL} .

In particular, \hat{R}^{HL} provides a uniform upper envelope for the Goldbach remainder but does not capture its pointwise asymptotic behaviour. \square

Remark (Stress–Test Regimes for Scale–Only Envelopes). Lemma A.2 shows that the Goldbach remainder $R^{HL}(2n; L)$ cannot be captured, even asymptotically, by any envelope depending only on the scale L up to a vanishing relative error. The obstruction is the intrinsic exponential interpolation of the remainder within cutoff cells, which is not controlled by scale alone.

Consequently, the framework identifies admissible regimes in which any uniform $O(L)$ or $O(\sqrt{2M})$ –type bound must necessarily be either non–tight or overly conservative. These regimes provide natural stress tests for scale–based remainder models.

A.3. Bounding Envelope Constant Per Envelope

Lemma A.3 (Bounding Envelope Constant Per Envelope).

Assume the analytic remainder bounding envelope is written in the form

$$\hat{R}^{HL}(2n; L) := \frac{c(2n; L)L}{\mathfrak{S}_{\text{head}}^{\mathcal{Q}, \mathcal{C}}(2n; L)},$$

where $c(2n; L)$ is chosen so that

$$R^{HL}(2n; L) \leq \frac{c(2n; L)L}{\mathfrak{S}_{\text{head}}^{\mathcal{Q}, \mathcal{C}}(2n; L)} \quad \text{for all admissible } L \geq p_{\min}(n) - 1.$$

Fix n , and let $p = p(L)$ be the cutoff prime determined by

$$\mathcal{Q}_p(n) \leq L < \mathcal{Q}_{p^+}(n),$$

where p^+ denotes the next cutoff prime after p in the cutoff order. (Equivalently, p is the largest cutoff prime with $\mathcal{Q}_p(n) \leq L$.) Then the envelope constant is given explicitly by

$$c(2n; L) = c\left(2n; \mathcal{Q}_p(n)\right) = 2 \prod_{\substack{q \in \mathbb{P}_{\text{eff}}(n) \\ \mathcal{Q}_q(n) > \mathcal{Q}_p(n)}} (q-2)^{\mathcal{Q}_p(n)/\mathcal{Q}_q(n)}.$$

In particular, $c(2n; L)$ is constant on each envelope interval $[\mathcal{Q}_p(n), \mathcal{Q}_{p^+}(n))$.

Proof. Fix n . On a cutoff envelope interval $[\mathcal{Q}_p(n), \mathcal{Q}_p^+(n))$, the index set $\{q : \mathcal{Q}_q(n) \leq L\}$ is unchanged, so the envelope normalization used to define $c(2n; L)$ is unchanged; hence $c(2n; L)$ is constant on this interval and equals its cutoff value $c(2n; \mathcal{Q}_p(n))$.

At the cutoff $L = \mathcal{Q}_p(n)$, the exact remainder in the product form is expressed as

$$R^{HL}(2n; \mathcal{Q}_p) = 2 \frac{\mathcal{Q}(2n; L)}{\mathfrak{S}_{\text{head}}^{\mathcal{Q}, \mathcal{C}}(2n; \mathcal{Q}_p)} \Xi(2n; \mathcal{Q}_p),$$

Using the envelope definition

$$c(2n; \mathcal{Q}_p) := \frac{R^{HL}(2n; \mathcal{Q}_p) \mathfrak{S}_{\text{head}}^{\mathcal{Q}, \mathcal{C}}(2n; \mathcal{Q}_p)}{\mathcal{Q}_p},$$

then

$$c(2n; \mathcal{Q}_p) = 2 \mathcal{Q}_p \mathfrak{S}_{\text{head}}^{\mathcal{Q}, \mathcal{C}}(2n; \mathcal{Q}_p) \left(\prod_{\substack{q \in \mathbb{P}_{\text{eff}}(n) \\ q > p}} (q-2)^{\mathcal{Q}_p / \mathcal{Q}_q} \right) / \left(\mathfrak{S}_{\text{head}}^{\mathcal{Q}, \mathcal{C}}(2n; \mathcal{Q}_p) \mathcal{Q}_p \right)$$

Simplifying,

$$c(2n; \mathcal{Q}_p) = 2 \prod_{\substack{q \in \mathbb{P}_{\text{eff}}(n) \\ q > p}} (q-2)^{\mathcal{Q}_p / \mathcal{Q}_q}.$$

Since $c(2n; L)$ is constant on $[\mathcal{Q}_p, \mathcal{Q}_p^+)$, this also gives $c(2n; L) = c(2n; \mathcal{Q}_p)$ for all L in the envelope, proving (A.25). \square

A.4. Monotonic Envelope

Lemma A.4 (Monotonic Envelope Constant for the Remainder Bounding Envelope).

Assume the analytic remainder bounding envelope is written in the form

$$\hat{R}^{HL}(2n; L) := \frac{c(2n; L)L}{\mathfrak{S}_{\text{head}}^{\mathcal{Q}, \mathcal{C}}(2n; L)},$$

where $c(2n; L)$ is chosen so that

$$R^{HL}(2n; L) \leq \frac{c(2n; L)L}{\mathfrak{S}_{\text{head}}^{\mathcal{Q}, \mathcal{C}}(2n; L)} \quad \text{for all admissible } L.$$

Then $c(2n; L)$ is monotone nonincreasing in L :

$$c(2n; L + \Delta L) \leq c(2n; L) \quad \text{for all } \Delta L \geq 0.$$

Moreover,

$$\lim_{L \rightarrow \infty} c(2n; L) = 2.$$

Proof. Fix n . Let the cutoff scales be

$$\mathcal{Q}_{p_0} < \mathcal{Q}_{p_1} < \mathcal{Q}_{p_2} < \dots,$$

where $p_0 < p_1 < p_2 < \dots$ are consecutive odd primes in the cutoff order, and \mathcal{Q}_p denotes the effective modulus at prime p (with base prime $p_{\min}(n) \geq 5$).

Case 1: L and $L + \Delta L$ lie in the same cutoff envelope. If

$$L, L + \Delta L \in [\mathcal{Q}_{p_0}, \mathcal{Q}_{p_1}),$$

then the set $\{p : \mathcal{Q}_p \leq L\}$ is unchanged, hence the envelope normalization defining $c(2n; L)$ is unchanged.

Therefore

$$c(2n; L + \Delta L) = c(2n; L).$$

At a cutoff scale $L = \mathcal{Q}_p(n)$, the envelope constant admits an explicit representation in terms of the geometric tail. By Lemma A.3;

$$c(2n; \mathcal{Q}_p) = 2 \prod_{q > p} (q-2)^{\mathcal{Q}_p / \mathcal{Q}_q}.$$

This expression is exact and characterizes the envelope constant associated with the cutoff prime p .

Taking logarithms yields

$$\log c(2n; \mathcal{Q}_p) = \log 2 + \mathcal{Q}_p \sum_{q > p} \frac{\log(q-2)}{\mathcal{Q}_q},$$

where the sum runs over the primes in the geometric tail. The superfactorial growth of \mathcal{Q}_q ensures absolute convergence of the series and makes the dependence on the cutoff p explicit.

Let $p_0 < p_1$ be consecutive cutoff primes, and enumerate the tail primes as

$$p_1 < p_2 < p_3 < \dots, p_i > p_0.$$

Then subtracting (A.38) at p_0 and p_1 yields

$$\begin{aligned} f(2n; \mathcal{Q}_{p_0}) &:= \log c(2n; \mathcal{Q}_{p_1}) - \log c(2n; \mathcal{Q}_{p_0}) \\ &= \mathcal{Q}_{p_0} \sum_{i=1}^{\infty} \frac{(p_1-1)\log(p_{i+1}-2) - (p_{i+1}-1)\log(p_i-2)}{(p_{i+1}-1)\mathcal{Q}_{p_i}}. \end{aligned}$$

All weights $\frac{\mathcal{Q}_{p_0}}{(p_{i+1}-1)\mathcal{Q}_{p_i}}$ are strictly positive, so the sign of each summand is determined by

$$g_i := (p_1 - 1)\log(p_{i+1} - 2) - (p_{i+1} - 1)\log(p_i - 2).$$

Sign of g_i . Since $p_{\min}(n) \geq 3$, we have $p_i \geq 3$ for all $i \geq 1$. Define,

$$\phi(x) := \frac{\log(x-2)}{x-1}.$$

If $x = 3$ then $\phi(x) = 0$.

Otherwise, a direct derivative computation shows

$$\phi'(x) = \frac{\frac{x-1}{x-2} - \log(x-2)}{(x-1)^2} < 0 \quad (x \geq 7),$$

and

$$\phi'(x) = 0 \quad (x = 5),$$

so ϕ is strictly decreasing on $[7, \infty)$. Hence, for all $i \geq 1$,

$$\frac{\log(p_{i+1} - 2)}{p_{i+1} - 1} < \frac{\log(p_i - 2)}{p_i - 1} \iff (p_i - 1)\log(p_{i+1} - 2) < (p_{i+1} - 1)\log(p_i - 2).$$

Since $p_1 \leq p_i$, we have $p_1 - 1 \leq p_i - 1$, and therefore

$$(p_1 - 1)\log(p_{i+1} - 2) \leq (p_i - 1)\log(p_{i+1} - 2) < (p_{i+1} - 1)\log(p_i - 2),$$

which implies $g_i < 0$ for all $i \geq 1$ where $p_i \geq 7$.

Thus every summand in (A.40) is strictly negative, and

$$f(2n; \mathcal{Q}_{p_0}) < 0, \text{ hence} \quad c(2n; \mathcal{Q}_{p_1}) < c(2n; \mathcal{Q}_{p_0}).$$

This proves strict decrease across each cutoff, and therefore monotone nonincrease for all $\Delta L \geq 0$.

Limit as $L \rightarrow \infty$. From (A.25) and the superfactorial growth of \mathcal{Q}_q , we have $\mathcal{Q}_p / \mathcal{Q}_q \rightarrow 0$ rapidly as $q \rightarrow \infty$, so each tail factor satisfies

$$\lim_{q \rightarrow \infty} (q-2)^{\mathcal{Q}_p / \mathcal{Q}_q} = 1.$$

Hence the infinite tail product converges to 1, and

$$\lim_{p \rightarrow \infty} c(2n; \mathcal{Q}_p) = 2.$$

Since $c(2n; L)$ is constant within each cutoff envelope, this implies

$$\lim_{L \rightarrow \infty} c(2n; L) = 2.$$

□

References

1. ^{a, b, c, d, e, f}Iwaniec H, Kowalski E (2004). *Analytic Number Theory*. Providence (RI): American Mathematical Society.
2. ^ABombieri E (1965). "On the Large Sieve." *Mathematika*. **12**:201–225. doi:[10.1112/S0025579300005313](https://doi.org/10.1112/S0025579300005313).
3. ^AVinogradov AI (1965). "The Density Hypothesis for Dirichlet L-Series." *Izv Akad Nauk SSSR Ser Mat [Proceedings of the USSR Academy of Sciences. Mathematical Series]*. **29**:903–934.
4. ^{a, b, c, d}Hardy GH, Littlewood JE (1923). "Some Problems of 'Partitio Numerorum', III: On the Expression of a Number as a Sum of Primes." *Acta Math*. **44**(1):1–70. doi:[10.1007/BF02403921](https://doi.org/10.1007/BF02403921).
5. ^{a, b, c, d, e, f, g}Riemers BC (2025). "Sieve-Theoretic Reformulation of Goldbach's Conjecture." Zenodo. <https://doi.org/10.5281/zenodo.18338749>.
6. ^{a, b}Harman G (2007). *Prime-Detecting Sieves*. Princeton (NJ): Princeton University Press.
7. ^AHardy GH, Wright EM (2008). *An Introduction to the Theory of Numbers*. 6th ed. Oxford: Oxford University Press.
8. ^AAbramowitz M, Stegun IA (1964). "Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables." National Bureau of Standards.
9. ^ACorless RM, Gonnet GH, Hare DEG, Jeffrey DJ, Knuth DE (1996). "On the Lambert W Function." *Adv Comput Math*. **5**(1):329–359. doi:[10.1007/BF02124750](https://doi.org/10.1007/BF02124750).
10. ^ALeCun Y, Chopra S, Hadsell R (2007). "A Tutorial on Energy-Based Learning." In: Bakir G, Hofmann T, Schölkopf B, Smola AJ, Taskar B, Vishwanathan SVN, editors. *Predicting Structured Data*. Cambridge (MA): MIT Press. p. 191.
11. ^ABottou L, Curtis FE, Nocedal J (2018). "Optimization Methods for Large-Scale Machine Learning." *SIAM Rev*. **60**(2):223–311.
12. ^ABottou L (2012). "Stochastic Gradient Descent Tricks." In: Montavon G, Orr GB, Müller KR, editors. *Neural Networks: Tricks of the Trade*. 2nd ed. Berlin: Springer. p. 421–436.
13. ^AHaffner P, LeCun Y (1991). "Graph Transformer Networks for Image Classification." In: *Advances in Neural Information Processing Systems*.

Declarations

Funding: No specific funding was received for this work.

Potential competing interests: No potential competing interests to declare.