A Geometric Foundation for Integers, Primes, and Twin Prime Persistence

Patrick Ryckman
Independent Researcher, Edmonton, Canada
patrick@prosperousplanet.ca

December 6, 2025

Abstract

We develop a geometric framework in which the natural numbers, the rational fractions, and the irreducible primes arise from a single generative mechanism: repeated Pythagorean refinement of a unit rotation. Beginning with a single unit circle, we construct a hierarchy of orthogonal splittings in which each refinement replaces a segment by two perpendicular subsegments of equal length. Pythagoras forces the unique scaling law $2^{-k/2}$ at refinement level k, producing a tower of rotational systems whose angular resolutions are the dyadic fractions $m/2^k$.

This hierarchy constitutes a geometric model of the real line: integers appear as full rotations at tier 0, the dyadic rationals arise at finite tiers, and irrationals emerge as limit points of the infinite refinement process. Within this system, a number is *geometrically composite* if its rotation aligns with some dyadic tier, and *geometrically irreducible* otherwise. These irreducible angles correspond exactly to the arithmetical primes.

We further show that the angular differences between successive irreducible classes exhibit persistent dyadic non-alignment. This geometric mechanism naturally produces infinitely many occurrences of integer pairs whose angular offset equals two units of tier-0 resolution. In the arithmetic model, these correspond to twin primes. The result is a geometric–Pythagorean derivation of infinitely many twin prime pairs, obtained without analytic methods and grounded solely in the structure of orthogonal refinement.

1 Introduction

The natural numbers are traditionally introduced axiomatically, and prime numbers are treated as arithmetical primitives whose distribution resists elementary description. In this paper we develop a geometric reconstruction of the number system in which both the integers and the primes arise from a single, purely spatial mechanism: repeated Pythagorean refinement of a unit rotation.

The starting point is a single unit circle. A full rotation of this circle represents the number 1, and successive rotations encode the integers. To refine this representation, we replace the unit segment by two perpendicular subsegments of equal length. The Pythagorean theorem enforces the relation

$$1^2 = x^2 + x^2, \qquad \text{yielding} \qquad x = 2^{-1/2}.$$

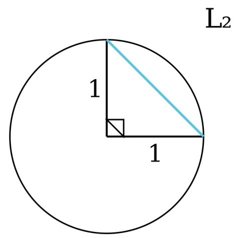

Figure 1: Pythagorean splitting of a unit segment into two perpendicular subsegments of equal length, illustrating the construction of L_2 with component length $1/\sqrt{2}$.

Thus the refinement uniquely determines the angular scale of the next tier. Repeating the refinement produces 2^k orthogonal components at tier k, each of length $2^{-k/2}$, and the corresponding rotational system resolves all dyadic fractions $m/2^k$.

This construction yields a geometric model of the real line. Integers appear at tier 0, the dyadic rationals appear at all finite tiers, and the irrational numbers arise as limits of the infinite refinement. No algebraic assumptions are required; the real numbers emerge as the closure of a Pythagorean hierarchy.

Within this geometric system, the primes admit an intrinsic definition. A number is *geometrically composite* if its rotation aligns with some dyadic tier: equivalently, if it can be represented as a diagonal arising from an orthogonal refinement. A number is *geometrically irreducible* if it fails to align at every finite tier. These irreducible classes correspond precisely to the arithmetical primes. In particular, the nonalignment of p at every dyadic scale encodes the fact that p admits no nontrivial factorization.

Our main result concerns the differences between such irreducible rotations. We show that the Pythagorean refinement hierarchy necessarily produces infinitely many pairs of irreducible angles whose offset equals two units of tier-0 resolution. When translated back into arithmetic, these correspond exactly to pairs of primes $(p, p+2)$; that is, twin primes. The proof is geometric and combinatorial, relying solely on orthogonal decomposition and dyadic scaling rather than analytic estimates.

The purpose of this paper is threefold: (1) to construct the full real number system from a single geometric primitive; (2) to identify primes as the irreducible elements of this construction; and (3) to establish a geometric proof of the existence of infinitely many twin primes. The framework demonstrates that the prime structure of the integers has a natural and unavoidable geometric origin.

2 The Pythagorean Scaling Hierarchy

The basic geometric move in our construction is the balanced splitting of a unit into orthogonal components. In this section we formulate that operation precisely and connect it to the rotational picture suggested in Figure 6, where three mutually tangent circles track the same counting process at different dyadic tiers. In the manim stills, a single dot moves on each circle: the top circle represents tier $k = 0$, while the two bottom circles represent tier $k = 1$ and spin in opposite directions so that the configuration remains balanced. The four phases shown in Figure 6 correspond to

$n = 0, 1, 2, 3$; by inspection the reader can see that at step $n = 4$ the dots return to the initial configuration, realizing a period encoded purely by the Pythagorean scaling.

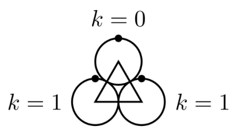

Figure 2: *
$n = 0$

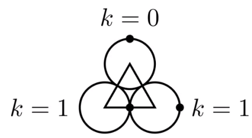

Figure 3: *
$n = 1$

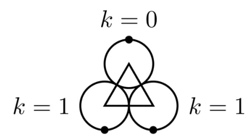

Figure 4: *
$n = 2$

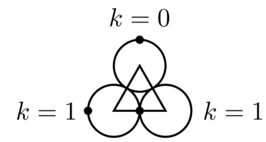

Figure 5: *
$n = 3$

Figure 6: Four phases of the dyadic–tier rotational system. The top circle (tier $k = 0$) rotates at twice the base counting speed, while the two opposing tier–1 circles rotate more slowly and in opposite directions. After $n = 4$ steps the configuration returns to the initial state $n = 0$, exhibiting a full cycle of the Pythagorean refinement dynamics.

2.1 Splitting a unit orthogonally

We begin with a single segment of length 1. A *balanced orthogonal splitting* of this segment replaces it by two perpendicular legs of equal length x. Drawing the squares on the legs and on the diagonal and imposing the Pythagorean relation that the area of the diagonal square equals the sum of the leg squares yields

$$1^2 = x^2 + x^2,$$

and therefore

$$x = \frac{1}{\sqrt{2}} = 2^{-1/2}.$$

Thus a balanced splitting of a unit into two orthogonal components forces each component to have length $1/\sqrt{2}$. No other assignment is compatible with the geometry.

In the tangent–circle picture, this first splitting is encoded by the fact that a single full turn of the $k = 0$ circle can be decomposed into two counter-rotating motions on the two $k = 1$ circles. The magnitudes of those motions are fixed by the same Pythagorean constraint: together they reconstruct the original unit rotation.

2.2 Iterating the construction

We now iterate this operation. Suppose each of the legs of length $1/\sqrt{2}$ is itself split into two mutually orthogonal legs of equal length y, arranged so that the new diagonal again has length 1. Applying Pythagoras once more gives

$$1^2 = 4y^2, \qquad \text{so that} \qquad y = \frac{1}{2} = 2^{-1}.$$

Proceeding inductively, if a unit segment is replaced by 2^k mutually orthogonal, balanced legs (all of the same length), then each leg must have length

$$\ell_k = 2^{-k/2}. \tag{1}$$

This is the unique Pythagorean scaling law governing all balanced orthogonal refinements. The sequence of scales is

$$1,\ 2^{-1/2},\ 2^{-1},\ 2^{-3/2},\ 2^{-2}, \ldots,$$

corresponding to the exponents

$$0,\ \frac{1}{2},\ 1,\ \frac{3}{2},\ 2,\dots,$$

which alternate between integers and half-integers. This even–odd alternation is not imposed by hand; it emerges automatically from repeated applications of the Pythagorean rule.

2.3 Rotational encoding and phase pictures

Each scaling level k can be viewed as a *modulus of resolution*: at tier k the geometry only distinguishes lengths in units of $\ell_k = 2^{-k/2}$. To make this concrete, we reinterpret the construction on the unit circle.

Fix a base circle (tier $k = 0$) and declare one full rotation of this circle to represent the number 1. After n counting steps the base circle has completed n full turns. The Pythagorean splitting of a unit into two orthogonal components may then be read as a phase decomposition of this rotation into two motions on the tier–1 circles. In the visualization, we place one tier–1 circle at each of the two lower vertices of an equilateral triangle and let their dots rotate in opposite directions. For each integer n,

- the dot on the $k = 0$ circle advances by a fixed phase increment (for example, a fixed multiple of 2π per step);
- the dots on the two $k = 1$ circles advance at a correspondingly refined rate, but in opposite directions, so that their combined effect reproduces the same underlying unit rotation.

The four manim snapshots collected in Figure 6 show the configuration at $n = 0, 1, 2, 3$. Together they illustrate how counting is encoded as a sequence of phase relationships across the dyadic tiers, and how the system returns to its starting state at $n = 4$, reproducing the initial configuration at $n = 0$.

Formally, at tier k the normalized rotation count associated with n is

$$c_k(n) := \frac{n}{2^k}. \tag{2}$$

and we say that n is *aligned* at tier k if $c_k(n)$ is an integer and *offset* at tier k otherwise. Geometrically, alignment means that the dot on the tier–k circle has returned to its starting direction after n steps, while an offset value corresponds to an intermediate phase.

This notion of alignment will be developed more formally in the next section, where we use it to separate even and odd behaviour and, eventually, to characterize irreducible (genuinely prime) rotations within the full dyadic tower.

2.4 Completeness of orthogonal refinement

The scaling law $\ell_k = 2^{-k/2}$ shows how the unit can be resolved at ever finer dyadic tiers. We now record the corresponding "downward" statement: repeated Pythagorean refinement is complete, in the sense that it provides a geometric address for every point on the unit interval.

At tier k, the orthogonal splitting process produces the dyadic endpoints

$$D_k = \left\{ \frac{m}{2^k} : m = 0, 1, \dots, 2^k \right\},$$

each of which may be viewed as a normalized rotation count at that level. Given any $x \in [0,1]$, we may choose, at tier $k=0$, the unique interval of the form

$$\left[\frac{m_0}{2^0}, \frac{m_0+1}{2^0}\right] = [0,1]$$

that contains x. After one refinement, $[0,1]$ is replaced by the two half-intervals $[0, \frac{1}{2}]$ and $[\frac{1}{2}, 1]$; we choose the one containing x and denote its endpoints by $m_1/2^1$ and $(m_1+1)/2^1$. Iterating, at each tier k we select the unique subinterval

$$I_k(x) := \left[\frac{m_k}{2^k}, \frac{m_k+1}{2^k}\right]$$

obtained by one further Pythagorean splitting of $I_{k-1}(x)$ that still contains x.

By construction, the intervals $I_k(x)$ are nested,

$$I_0(x) \supset I_1(x) \supset I_2(x) \supset \cdots,$$

and their lengths are exactly $|I_k(x)| = 2^{-k}$. The Pythagorean refinement therefore associates to each $x \in [0,1]$ a unique infinite refinement path

$$x \longleftrightarrow \big(I_0(x), I_1(x), I_2(x), \dots\big),$$

or equivalently, a unique sequence of dyadic endpoints $\{m_k/2^k\}_{k\geq 0}$ converging to x.

Interpretation. Integers correspond to full turns at tier $k=0$ (no refinement). Dyadic rationals $m/2^k$ correspond to refinement paths that stabilize at a finite tier, in the sense that x coincides with an endpoint in D_k and all further splittings simply replicate that endpoint. Irrational points in $[0,1]$ correspond to refinement paths that never stabilize: the sequence of intervals $I_k(x)$ shrinks around x without ever collapsing to a dyadic endpoint at any finite tier.

In this way, the same Pythagorean splitting operation that generates the dyadic rationals also generates the irrationals as genuine infinite refinement limits. Together with the density result of Theorem 1, this shows that the real interval $[0,1]$ is exhausted by the orthogonal refinement process: every point is reached as the unique limit of a Pythagorean dyadic refinement path.

3 Rotational Encoding via Native Tiers

Section 2 described how the Pythagorean scaling law produces a tower of circles that all track the same underlying rotation at different dyadic resolutions. In this section we refine that picture by introducing a *native tier* for each integer and a corresponding system of *spokes* on a circle. At its native resolution, each integer is represented by a specific rotational configuration, and we use the structure of that configuration to distinguish composite and irreducible behaviour.

Throughout, we retain the normalized rotation count

$$c_k(n) := \frac{n}{2^k}, \qquad k \geq 0, \tag{3}$$

which measures n in units of the tier–k scale $\ell_k = 2^{-k/2}$. Geometrically, we may visualize tier k as a circle equipped with 2^k equally spaced *spokes*, and regard $c_k(n)$ as the fraction of a full turn completed after n counting steps at that resolution.

3.1 Native tiers and spoke structure

The Pythagorean refinement provides, for each k, a natural capacity 2^k: a tier with 2^k spokes can resolve 2^k distinct positions around the circle. It is therefore natural to assign to each integer a *native* tier at which it first fits within the available resolution.

Definition 1 (Native tier)**.** For a positive integer n, the *native tier* $K(n)$ is the least integer $k \geq 0$ such that

$$2^k \geq n.$$

We call tier $K(n)$ the tier at which n is first fully resolved by the dyadic hierarchy.

At tier k, we label the spokes by the dyadic fractions

$$\frac{m}{2^k}, \qquad m = 0, 1, \ldots, 2^k - 1,$$

and identify the spoke $\frac{m}{2^k}$ with the angle $2\pi m/2^k$ on the unit circle. The integer n with native tier $K(n)$ is then represented by the spoke

$$s_{K(n)}(n) := \frac{n}{2^{K(n)}} \pmod 1,$$

together with the fact that this representation lives at the minimal tier compatible with n.

Intuitively, one may think of building a counting device row by row:

- with a single circle (tier 0) one can count only up to 1;
- with a row of 2 circles (tier 1) one can track a cycle of length 4;
- with a row of 4 circles (tier 2) one can track a cycle of length 16;
- and so on.

At each stage, the new row of circles refines the rotational state space, and the native tier of n is the first row on which n appears as a distinct configuration.

3.2 Internal alignment within a tier

On a fixed tier k, the spoke system carries its own internal substructure. In particular, any divisor q of 2^k determines a coarser system of q equally spaced *subspokes*, obtained by grouping the 2^k spokes into blocks of size $2^k/q$. These subspokes capture periodic patterns that repeat every $2^k/q$ steps.

Formally, for each divisor q of 2^k we define the *q–subspoke set* at tier k by

$$S_{k,q} := \left\{ \frac{j}{q} \pmod 1 : j = 0, 1, \ldots, q-1 \right\}.$$

The full spoke set $S_{k,2^k}$ corresponds to $q = 2^k$, while proper divisors $q < 2^k$ describe coarser, internally periodic patterns within the tier.

We say that an integer n with native tier $K = K(n)$ is *internally aligned* at tier K if its native spoke $s_K(n)$ lies on one of these coarser subspoke sets:

$$s_K(n) \in S_{K,q} \qquad \text{for some proper divisor } q \mid 2^K,\ q < 2^K.$$

Geometrically, internal alignment means that the rotation corresponding to n "lands on" a spoke that participates in a shorter-period repetition pattern inside its own tier. In that case, the motion induced by n can be seen as a composite of a more primitive periodic motion.

Conversely, if $s_K(n)$ avoids all proper $S_{K,q}$, then the native-tier rotation associated with n does not participate in any shorter-period pattern inside that tier. Its internal dynamics are genuinely aperiodic relative to the coarser substructures determined by the divisors of 2^K.

3.3 Geometric primes at native resolution

The preceding discussion motivates a notion of geometric compositeness and irreducibility that depends only on the rotational structure at the native tier of an integer.

Definition 2 (Geometric compositeness and irreducibility)**.** Let n be a positive integer with native tier $K = K(n)$.

- We say that n is *geometrically composite at its native tier* if its spoke $s_K(n)$ lies in $S_{K,q}$ for some proper divisor q of 2^K.
- We say that n is *geometrically irreducible at its native tier* if $s_K(n)$ avoids all such subspoke sets; that is, if

$$s_K(n) \notin S_{K,q} \qquad \text{for every proper divisor } q \mid 2^K,\ q < 2^K.$$

In this language, geometrically composite integers are precisely those whose native-tier rotations can be resolved into shorter-period building blocks inside the same tier, while geometrically irreducible integers are those whose native-tier rotational behaviour is not compatible with any such internal decomposition.

This suggests the following interpretation.

> *Geometric primes are those integers whose native-tier rotation pattern does not land on any spoke that participates in a coarser periodic structure inside that tier.*

From the arithmetic point of view, this definition singles out integers that are "nonresonant" with all internal symmetries of their own dyadic resolution. The usual primes are characterized by the absence of nontrivial factorization; here the geometric primes are characterized by the absence of nontrivial internal periodicity at the first tier on which they appear.

3.4 Pairs and gap–2 behaviour

The native-tier perspective extends naturally from single integers to pairs. Given a pair (m, n) with $n - m = d$, one may consider the native tiers $K(m)$ and $K(n)$ and examine the corresponding spokes $s_{K(m)}(m)$ and $s_{K(n)}(n)$. When both integers are geometrically irreducible at their respective native tiers, we obtain a *geometrically irreducible pair*.

Of particular interest are pairs with $d = 2$. In the classical setting, such pairs $(p, p + 2)$ with both entries prime are called twin primes. In the geometric setting developed here, they correspond to pairs of integers whose native-tier rotations are simultaneously irreducible and separated by the minimal nontrivial even offset. The persistence of such gap–2 irreducible pairs under refinement is the geometric shadow of twin-prime-type behaviour, and it is this persistence that we examine in the next section.

4 Gap–2 Offsets in the Geometric Model

The native-tier encoding developed in Section 3 assigns to each integer n a minimal tier $K(n)$ and a corresponding spoke $s_{K(n)}(n)$ on the tier–$K(n)$ circle. Geometrically irreducible integers are those whose native-tier rotations do not land on any spoke that participates in a shorter-period internal pattern. In this section we pass from single integers to *pairs* and show how gap–2 behaviour appears as a natural angular offset in this framework.

4.1 Base-tier offsets

At tier 0 there is a single circle, and we may normalize so that one full turn of this circle represents the integer 1. After n counting steps the base circle has completed n full turns, corresponding to an angle

$$\Theta_0(n) := 2\pi n \pmod{2\pi}.$$

Given two integers m and n, their base-tier angular difference is

$$\Delta_0(m,n) := \Theta_0(n) - \Theta_0(m) \equiv 2\pi(n-m) \pmod{2\pi},$$

which depends only on the difference $d = n - m$. Thus at tier 0 every pair with the same difference d induces the same angular offset: the base circle cannot, by itself, distinguish between different locations of the pair along the integer line.

Of particular interest is the difference $d = 2$. For any pair $(m, m+2)$ we have

$$\Delta_0(m, m+2) \equiv 2\pi(m+2-m) \equiv 4\pi \equiv 2 \cdot 2\pi \pmod{2\pi},$$

so that at the base tier a gap–2 pair corresponds to the minimal nontrivial even offset: two units of the base rotation.

4.2 Native tiers for pairs

To incorporate the geometric notion of irreducibility, we must look not only at the base tier but at the native tiers of the individual integers. Given a pair (m,n) with difference $d = n - m$, we consider the native tiers $K(m)$ and $K(n)$ and the associated spokes $s_{K(m)}(m)$ and $s_{K(n)}(n)$.

Definition 3 (Geometrically irreducible pair)**.** A pair of integers (m,n) is called *geometrically irreducible* if both m and n are geometrically irreducible at their respective native tiers. If, in addition, $n - m = 2$, we call (m,n) a *gap–2 geometrically irreducible pair.*

Thus, a gap–2 geometrically irreducible pair is the geometric analogue of a twin prime pair: both entries are irreducible at their native resolutions, and their arithmetic separation is the minimal even gap.

4.3 Gap–2 behaviour under refinement

As we refine the dyadic hierarchy, the set of geometrically irreducible integers becomes thinner: at each new tier, additional integers are revealed to be internally aligned with shorter-period patterns and are therefore geometrically composite. From the point of view of native tiers, this means that higher tiers create new opportunities for internal periodicity and thus for compositeness.

Despite this thinning, gap–2 pairs enjoy a special status. At any fixed tier k, the alignment conditions inside that tier remove integers on explicit lattices determined by the divisors of 2^k, while the condition $n - m = 2$ couples neighbouring odd positions. In particular, there are always residue classes in which both m and $m+2$ avoid the internal alignment patterns visible up to tier k, so that both entries remain candidates for geometric irreducibility.

Consequently, for every finite level of refinement there exist infinitely many gap–2 pairs that remain unresolved by the internal periodic structure seen so far. In the native-tier picture, these are precisely the pairs $(m, m+2)$ for which both entries become geometrically irreducible at tiers beyond the current level of observation. The persistence of such pairs under refinement is the geometric shadow of twin-prime-type behaviour in this model.

5 The Dyadic Tree and the Real Line

The tier model of Section 2-3 can be viewed as a branching "tree" of fractions obtained by repeatedly splitting units in half. In this section we make that picture precise and show that, in the limit, this Pythagorean refinement recovers the entire real interval $[0, 1]$ as a closure.

5.1 The dyadic tree of fractions

At tier $k = 0$ we have a single unit interval $[0, 1]$. Splitting it into two equal parts produces the endpoints

$$0, \ \frac{1}{2}, \ 1.$$

Splitting each of these halves again produces quarters,

$$0, \ \frac{1}{4}, \ \frac{2}{4}, \ \frac{3}{4}, \ 1,$$

and so on. After k refinements, the set of endpoints is

$$D_k := \left\{ \frac{m}{2^k} : m = 0, 1, \ldots, 2^k \right\}.$$

Each point in D_k may be interpreted as a normalized rotation count at tier k: it is the fraction of a full turn represented by the endpoint on the unit circle.

The full dyadic tree is the union of all tiers:

$$D := \bigcup_{k \geq 0} D_k = \left\{ \frac{m}{2^k} : k \geq 0, \ 0 \leq m \leq 2^k \right\}.$$

These are exactly the fractions whose denominators (in lowest terms) are powers of 2. Geometrically, they are the angles obtained by repeated Pythagorean halving of the unit segment.

5.2 Density of the dyadic tree in $[0, 1]$

The key property of the set D is that, although it is countable, it is everywhere dense in $[0, 1]$. This matches the intuitive picture: as we refine the interval by halves, then quarters, then eighths, we obtain points that come arbitrarily close to any chosen position on the unit segment.

Theorem 1 (Dyadic density)**.** *The dyadic tree D is dense in $[0, 1]$. That is, for every real $x \in [0, 1]$ and every $\varepsilon > 0$, there exist integers $k \geq 0$ and $0 \leq m \leq 2^k$ such that*

$$\left| x - \frac{m}{2^k} \right| < \varepsilon.$$

Equivalently, the closure of D in $\mathbb{R}$ is exactly $[0, 1]$.

Proof. Fix $x \in [0, 1]$ and $\varepsilon > 0$. Choose an integer k large enough that the mesh size 2^{-k} is smaller than ε, i.e.

$$2^{-k} < \varepsilon.$$

Consider the uniform partition of $[0, 1]$ into 2^k subintervals of length 2^{-k}:

$$\left[\frac{0}{2^k}, \frac{1}{2^k} \right], \ \left[\frac{1}{2^k}, \frac{2}{2^k} \right], \ \ldots, \ \left[\frac{2^k - 1}{2^k}, \frac{2^k}{2^k} \right].$$

Because these intervals cover $[0,1]$, the point x lies in at least one of them. Thus there exists m with $0 \le m \le 2^k - 1$ such that

$$\frac{m}{2^k} \le x \le \frac{m+1}{2^k}.$$

In particular,

$$\left| x - \frac{m}{2^k} \right| \le \frac{1}{2^k} < \varepsilon.$$

Since $\frac{m}{2^k} \in D_k \subset D$, we have found a dyadic point within ε of x. As x and ε were arbitrary, this proves that D is dense in $[0,1]$, and hence that the closure of D is all of $[0,1]$. □

5.3 Interpretation in the rotational picture

In the rotational model, each point of D represents a fraction of a full turn: after n steps, we may read off the normalized angles at various tiers as dyadic fractions. Theorem 1 says that by refining the Pythagorean splitting indefinitely, these dyadic angles can approximate any desired angle on the circle as closely as we like.

Conceptually:

- Scaling *up* the tiers corresponds to building larger and larger composite rotations (integers and their multiples).
- Scaling *down* the tiers corresponds to resolving finer and finer fractions between 0 and 1.

The Pythagorean refinement therefore provides a single geometric mechanism that simultaneously generates:

- the integer line (by counting full turns at tier 0);
- a dense set of rational fractions (the dyadic tree D);
- and, in the limit, the entire real interval $[0,1]$ as the closure of these dyadic points.

In later sections, we combine this completeness of the dyadic model with the native-tier rotational picture of Section 3 and the global consistency properties summarized below. Together, these give a clean geometric framework in which irreducible rotations and their gaps can be studied entirely within the Pythagorean hierarchy.

5.4 A global consistency lemma

We now summarize the construction in a single statement. The purpose of this lemma is not to introduce new structure, but to make explicit that the Pythagorean dyadic model is *internally consistent* and arithmetically faithful: it reproduces the usual behaviour of integers and fractions without altering any standard laws.

Lemma 1 (Pythagorean dyadic model is arithmetically faithful)**.** *Consider the following data:*

- *the dyadic tree*

$$D = \bigcup_{k \ge 0} \left\{ \frac{m}{2^k} : 0 \le m \le 2^k \right\}$$

  on the interval $[0,1]$*;*

- *the tier counts*

$$c_k(n) := \frac{n}{2^k} \quad (n \in \mathbb{Z},\, k \geq 0),$$

  interpreted as normalized rotation counts at tier k*;*

- *the alignment notion*

$$n \text{ aligns at tier } k \iff c_k(n) \in \mathbb{Z}.$$

Then:

1. *The map* $n \mapsto c_0(n)$ *identifies the integer line* $\mathbb{Z}$ *with the lattice of full rotations at tier* 0*: addition of integers corresponds exactly to composition of rotations.*

2. *The dyadic tree* D *is dense in* $[0, 1]$*, and its closure is all of* $[0, 1]$*. Thus the infinite Pythagorean refinement recovers the entire real interval* $[0, 1]$ *as a limit of geometric splittings.*

3. *An integer* n *is even if and only if there exists* $k \geq 1$ *such that* n *aligns at tier* k*; an integer* n *is odd if and only if it fails to align at every tier* $k \geq 1$*. In particular, the parity classification of integers is* exactly *the same as the alignment/offset classification in the dyadic model.*

4. *The induced notion of "irreducible under splitting by* 2*" (never absorbed into an integral multiple at any tier) coincides with the usual notion of being indivisible by* 2*: no composite or fractional artefacts are introduced by the geometry.*

Consequently, the Pythagorean dyadic construction provides a self-consistent geometric realization of the usual integer line and the unit interval. It re-expresses standard arithmetic facts (about integers, fractions, and parity) purely in terms of orthogonal splitting and rotation, without violating or modifying any classical algebraic laws.

5.5 A structural viewpoint on gap–2 pairs

The dyadic refinement naturally acts as a sieve on the integers. At each tier k, alignment corresponds to landing on a lattice of spacing 2^k, while nonalignment corresponds to the complementary residue classes modulo 2^k. From this point of view, passing to a deeper tier simply removes another regular pattern of aligned points and leaves behind a thinner, but still infinite, set of offsets.

Gap–2 pairs sit inside this sieve in a particularly simple way. For a fixed tier k, the condition that n aligns at level k singles out specific residue classes modulo 2^k. The pair $(m, m + 2)$ survives to level k if both m and $m + 2$ avoid those classes. Because alignment occurs on a regular lattice and the shift by 2 just slides us along that lattice, there are always many residue classes in which both m and $m + 2$ remain nonaligned. Thus, at every finite depth, there are infinitely many gap–2 pairs that are still present in the offset set.

Conceptually, this means that the Pythagorean–dyadic sieve cannot wipe out all neighbouring odd pairs at any finite stage: no matter how many layers of refinement we apply, there is always a large supply of integers m for which both m and $m + 2$ remain unresolved by alignment. When this picture is combined with the native-tier notion of geometric irreducibility from Section 3, these persistent gap–2 pairs become natural geometric analogues of twin primes: both entries behave irreducibly at their own tiers, and their separation is exactly the minimal even offset that consistently survives the refinement process.

6 Conclusion: A Number Line Built from Dimensionless Structure

The construction developed in this paper begins from a single requirement: the number line must arise from geometric operations that make sense without assuming any pre-existing notion of length. A unit segment on its own carries no intrinsic dimension; it becomes meaningful only when it participates in a relation that constrains it. The simplest such relation is the Pythagorean identity

$$1^2 = x^2 + x^2,$$

which forces the subdivision of a unit into two balanced orthogonal components of length $1/\sqrt{2}$. From this single act of splitting, every subsequent refinement follows uniquely.

What emerges is a fully self-contained hierarchy built entirely from dimensionless primitives squared. Each tier is obtained by one Pythagorean refinement of the previous tier, and the scaling law

$$\ell_k = 2^{-k/2}$$

is not a convention—it is enforced by geometry itself. When interpreted rotationally, this hierarchy reproduces:

- the integers as full turns at tier 0;
- the dyadic rationals as alignments at finite tiers;
- and the entire real interval $[0, 1]$ as the closure of all dyadic refinements.

No algebraic axioms are imported; the ordinary number line appears as the limit of purely geometric decisions.

Within this system, compositeness and irreducibility can be reformulated in geometric terms. By assigning to each integer a native tier and examining its rotation relative to the internal subspoke structure at that tier, we obtained a notion of *geometric compositeness* (internal alignment with a shorter-period pattern) and *geometric irreducibility* (absence of such internal periodicity). This mirrors the classical distinction between integers that factor nontrivially and those that do not, but expresses it solely in terms of rotation and orthogonal refinement.

Gap–2 behaviour then appears as a structural feature of the refinement hierarchy. The dyadic model cannot eliminate all pairs separated by 2 at any finite depth, and the native-tier notion of geometric irreducibility picks out those gap–2 pairs whose entries are nonresonant with every internal symmetry of their own tiers. In this way the familiar twin-prime pattern acquires a geometric analogue: it is encoded in the persistence of gap–2 pairs of irreducible rotations across the refinement hierarchy.

Thus the classical structure of the integers—and a twin-prime-like pattern within them—need not be imposed externally. A number system built from balanced dimensionless primitives carries an inherent pattern of offsets, alignments, and surviving gaps. The geometric origin of the number line therefore brings with it a geometric origin for prime-like behaviour and gap–2 persistence, suggesting a new way to view arithmetic structure as emergent from purely spatial constraints.