Sparse Coefficient Spectra of Pell Representation Polynomials: Growth, Block Structure, and Failure of Halving Laws

HYDRA Research Programme
Collective project attribution; see provenance disclosure

6 November 2026

This manuscript originates from the HYDRA research programme and was developed under human direction with substantial AI contributions to mathematical discovery, proof writing, computational verification, and editing. The research-programme credit on the title page is a collective project attribution, not a legal authorship determination; platform-specific formal author and corresponding-author records are supplied as separate, Owner-approved submission metadata. AI mathematical reviews were completed in Track B (Kimi Work) and B-2 (GPT subagents), with explicitly documented limits on reviewer independence and procedural blindness for the central claims. No independent human peer review or machine-checked formal proof certification is claimed. The prior-art search was bounded, novelty is not established, and no priority claim is made.

Abstract

We study the set of distinct coefficient values of the polynomials $G_m(x) = \prod_{j=2}^{m+1}(1 + x^{P_j} + x^{2P_j})$, where P_j are the Pell numbers. A coefficient-pair recurrence identifies the spectra at consecutive odd and even lengths. The resulting spectrum V_k admits an exact representation by products of blocks drawn from levels of the Calkin–Wilf tree, with additive level cost. We give the bound $a_k \leq 1 + A(5/3)^k$, where $a_k = |V_k|$ includes zero and $A = 95367431640625/14169773449216 < 7$. For every $\lambda > \varphi$, we also obtain $a_k = O(\lambda^k)$, and deduce density zero within the natural coefficient-value interval. An atomic factorization argument shows that the upper exponential growth rates of the spectrum, primitive cost layers, first block appearances, and block levels coincide in a number $\rho \in [2/\varphi, \varphi]$. We identify the block sets with the denominator sets of the Kepler tree, using an explicit level-preserving transformation. Finally, an exact finite certificate for 3847 disproves four precisely defined internal halving properties. Neither the existence of an ordinary growth limit nor the value of ρ is established.

1. Introduction

The number of nonzero coefficients, the maximum coefficient, and the number of distinct coefficient values are different statistics of a representation polynomial. This paper concerns the third statistic, counting zero once when it occurs within the polynomial's degree interval. We derive its structure directly from shifted copies of Pell coefficient sequences.

The positive coprime-pair tree used below is classical. In particular, our maps $L(p, q) = (p, p + q)$ and $R(p, q) = (p + q, q)$ are the usual Calkin–Wilf operations. Kepler denominator levels and their elementary Fibonacci bounds occur in Mathar [M]; OEIS A086592 [O] records

a levelwise relation to numerator-plus-denominator sequences. We use these objects as structural tools, without claiming their discovery. The mathematical content under consideration here is their connection to the entire coefficient-value spectrum, the product-cost counting bounds, and the common upper growth exponent.

The related public Pell preprint [P] concerns maximal coefficients, their locations, and integer fibres in the silver ratio. The present scope is distinct coefficient values. Our literature comparison is bounded; priority for the spectrum results remains unresolved. The halving properties in Section 7 are internal propositions of the HYDRA investigation, not identified historical open problems.

All general arguments are separated from finite computational evidence. Some historical C5/C6 reviews involved an AI reviewer that had also contributed to their adjudication and to the exponent and witness arguments. Subsequent Track B review by Kimi Work accepted the arguments of the Mathar–Kepler bridge, exponent identity, cost criterion, and finite witness within its documented independence scope; its review of the foundational Sections 2–3 and 5 remained corroborative. B-2 reviewed those foundational sections mathematically, but its separate subagent conversations did not have enforced filesystem/tool isolation or independently attested model-family independence. These are AI proof reviews with explicit limits, not a claim of fully independent certification.

2. Pell polynomials and the coefficient-pair recurrence

Let $P_0 = 0$, $P_1 = 1$, and $P_{j+2} = 2P_{j+1} + P_j$. For $m \geq 1$, define

$$G_m(x) = \prod_{j=2}^{m+1} (1 + x^{P_j} + x^{2P_j}), \qquad f_m(n) = [x^n]G_m(x).$$

Extend f_m by zero outside its degree interval. Put $C_m = P_{m+1}$, $A_m = P_{m+2}$, and $d_m = A_m + C_m - 3$. The degree formula follows by induction, starting with $d_1 = 4$ and adding $2P_{m+2}$ at the next multiplication. Palindromicity gives $f_m(n) = f_m(d_m - n)$. Define full sets

$$\widehat{V}_m = \{f_m(n) : 0 \leq n \leq d_m\}, \quad \widehat{K}_m = \{(f_m(t), f_m(t + C_m)) : 0 \leq t \leq d_m - C_m\}.$$

Pairs are ordered; these are sets, not multisets. Write $\sigma(a, b) = a + b$ and $\Delta(U) = \{(u, u) : u \in U\}$.

Lemma 2.1 (one-step values). For $m \geq 2$,

$$\widehat{V}_{m+1} = \widehat{V}_m \cup \sigma(\widehat{K}_{m-1}).$$

Proof. Multiplication by $1 + x^{A_m} + x^{2A_m}$ creates three copies of f_m. Since $d_m < 2A_m$, at most two copies overlap. Every old value occurs in the initial unmodified part: reflect its index if necessary to obtain $n \leq d_m/2 < A_m$. The three support intervals leave no gap, since $d_m \geq A_m - 1$.

The overlapping values are $f_m(t) + f_m(t + A_m)$ for $0 \leq t \leq C_m - 3$; the other overlap has the same values. With $g = f_{m-1}$, we have $f_m(t) = g(t) + g(t - C_m) + g(t - 2C_m)$. Since $A_m = 2C_m + P_m$ and $\deg g = C_m + P_m - 3$, the indicated range gives $f_m(t) = g(t)$ and $f_m(t + A_m) = g(t + P_m)$. This is precisely the full range of $\widehat{K}_{m-1}$. □

Lemma 2.2 (central band). For every $m \geq 1$, the values of f_m on $[C_m - 2, A_m - 1]$ form $\widehat{V}_m$.

Proof. For $m = 1$, the band is $[0, 4]$. For $m = 2$, the band $[3, 11]$ contains both 0 and 1, the only values. Assume the result at $m - 2$, with $m \geq 3$. Set $h = f_{m-2}$, $w = P_m$, $v = P_{m-1}$, $C = 2w + v$, and $A = 2C + w$. Then $\deg h = w + v - 3$. The nine shifts of h in f_m are

$$0, w, 2w, C, C + w, C + 2w, 2C, 2C + w, 2C + 2w.$$

The successive gaps are w, w, v, w, w, v, w, w. The sum of two successive gaps is at least $w + v = \deg h + 3$, excluding triple overlap. Since $v \geq 2$, each gap is at most $\deg h + 1$, so no uncovered support gap creates an additional value.

Every coefficient is therefore a single value of h, a sum $h(t)+h(t+w)$ with $0 \leq t \leq v-3$, or a sum $h(t) + h(t + v)$ with $0 \leq t \leq w - 3$. By induction every single value, including zero, has a representative $t \in [v - 2, w - 1]$. At $n = C + w + t$, both neighbouring contributions vanish, since $t - w < 0$ and $t + w > \deg h$. Thus this single value is reproduced inside the required band. The two sum types are reproduced at $n = C + w + t$ and $n = C + t$, respectively, using the corresponding central adjacent copies. All these indices lie in $[C - 2, A - 1]$. The opposite inclusion is definitional. □

Lemma 2.3 (full pair recurrence). For $m \geq 3$,

$$\widehat{K}_m = L(\widehat{K}_{m-2}) \cup R(\widehat{K}_{m-2}) \cup \Delta(\widehat{V}_m).$$

Proof. Keep h, w, v, C from Lemma 2.2 and let $g = f_{m-1}$, $f = f_m$. Then

$$g(t) = h(t) + h(t - w) + h(t - 2w), \qquad f(t) = g(t) + g(t - C) + g(t - 2C).$$

The degrees are $w + v - 3, C + w - 3, 3C + w - 3$, respectively. On $0 \leq t \leq w - 3$, direct substitution gives

$$(f(t), f(t + C)) = (h(t), h(t) + h(t + v)).$$

These are exactly $L(\widehat{K}_{m-2})$. On $w - 2 \leq t \leq 2C - 1$, subtraction gives

$$f(t + C) - f(t) = g(t + C) - g(t - 2C) = 0.$$

These diagonal values cover $\widehat{V}_m$. Indeed, the part $[C-2, 2C-1]$ of the central band is already here, while reflection of $[2C, A - 1]$ about degree $3C + w - 3$ lies in $[C - 2, C + w - 3]$, also within this diagonal range. Finally, for $t = 2C + z$, $0 \leq z \leq w - 3$, substitution gives

$$(f(t), f(t + C)) = (h(z) + h(z + v), h(z + v)),$$

exactly $R(\widehat{K}_{m-2})$. The three ranges partition the full pair-index interval. □

Theorem 2.4 (parity and semantic binding). For every $k \geq 0$,

$$\widehat{V}_{2k+1} = \widehat{V}_{2k+2}, \qquad \widehat{K}_{2k+1} = \widehat{K}_{2k+2}.$$

Consequently, putting $V_k = \widehat{V}_{2k+1}$ and $K_k = \widehat{K}_{2k+1}$, we have

$$V_0 = \{0, 1\}, \quad K_0 = \{(0, 0), (1, 1)\},$$

$$V_{k+1} = V_k \cup \sigma(K_k), \qquad K_{k+1} = L(K_k) \cup R(K_k) \cup \Delta(V_{k+1}). \tag{SB}$$

Proof. The nonzero positions of G_2 are $0, 2, 4, 5, 7, 9, 10, 12, 14$, all with coefficient 1. Thus the asserted value and pair equalities at lengths 1 and 2 hold directly. For even $m \geq 4$, assume the pair equality at $m - 3, m - 2$. Lemma 2.1 gives

$$\widehat{V}_{m-1} = \widehat{V}_{m-2} \cup \sigma(\widehat{K}_{m-3}), \qquad \widehat{V}_m = \widehat{V}_{m-1} \cup \sigma(\widehat{K}_{m-2}).$$

The latter union adds nothing new. Lemma 2.3 at $m - 1, m$ now gives the pair equality. This simultaneous induction proves both statements. Applying Lemmas 2.1 and 2.3 with the resulting parity identification gives (SB). In particular, K_k is the set of coefficient pairs at separation P_{2k+2} in G_{2k+1}, not merely a syntactic construction. □

Lemma 2.5 (coordinates and maximum). The coordinates of every pair in K_k belong to V_k, $\Delta(V_k) \subseteq K_k$, and $\max V_k = 2^k$.

Proof. Coordinate containment holds at $k = 0$. If it holds at k, then the coordinates after L or R are $a, b, a + b$; the first two lie in $V_k \subseteq V_{k+1}$, and the third in $\sigma(K_k) \subseteq V_{k+1}$. The inserted diagonal also has coordinates in V_{k+1}. Diagonal containment is immediate from (SB) and the base. If $M = \max V_k$, all sums from K_k are at most $2M$, while $(M, M) \in K_k$ attains $2M$. Since $M \geq 1$, $\max V_{k+1} = 2M$, giving the formula. □

3. Recursive pairs and multiplicative blocks

Define $T_0 = \{(1,1)\}$, $T_{j+1} = L(T_j) \cup R(T_j)$, and, for $c \geq 1$,

$$W_c = \{p+q : (p,q) \in T_{c-1}\}, \qquad w_c = |W_c|, \quad a_k = |V_k|.$$

Every pair is positive and coprime. Subtracting the smaller coordinate from the larger recovers its unique parent unless it is $(1,1)$. This also proves that the 2^j words of length j give distinct pairs.

Theorem 3.1 (genealogy and products). For all $k \geq 0$,

$$K_k = \bigcup_{j=0}^{k} \{(up, uq) : u \in V_{k-j}, (p,q) \in T_j\},$$

$$V_k \setminus \{0\} = \left\{ \prod_{i=1}^{r} g_i : g_i \in W_{c_i},\ c_i \geq 1,\ \sum_i c_i \leq k,\ r \geq 0 \right\}. \tag{PR}$$

The empty product is 1 with cost zero; decompositions need not be unique.

Proof. Genealogy follows by induction from (SB). The root term is the inserted diagonal. Homogeneity of L, R carries all earlier scaled layers to the next layer. The zero scale correctly accounts for $(0,0)$.

For necessity in (PR), induct on k. Old values retain their decomposition. Every new positive sum from K_{k-1} is $u(p+q)$, where $u \in V_{k-1-j}$, $(p,q) \in T_j$, and $u > 0$. The cost of u is at most $k-1-j$ by induction and $p+q \in W_{j+1}$ adds cost $j+1$. Conversely, if $u \in V_s$ and $g \in W_c$, start with $(u,u) \in K_s$, follow a word of length $c-1$, then sum. This gives $ug \in V_{s+c}$. Iterate from 1; monotonicity of V_k handles unused cost. The base $k = 0$ is $\{1\}$. □

Lemma 3.2 (block range). With $F_0 = 0, F_1 = 1, F_{j+2} = F_{j+1} + F_j$, every $g \in W_c$ satisfies $c+1 \leq g \leq F_{c+2}$.

Proof. Starting with sum 2, every step increases the sum by at least 1. For the upper bound, sort a pair as $x \leq y$. At depth j, induction gives $x \leq F_{j+1}, y \leq F_{j+2}$. A sorted child has smaller coordinate at most y and larger coordinate $x+y$, so the induction continues. At depth $c-1$ the sum is at most $F_c + F_{c+1}$. □

4. The Kepler connection

Let $\mathcal{F}_0 = \{1/1\}$, $\mathcal{F}_1 = \{1/2\}$, and let every i/j at level $c \geq 1$ have children $i/(i+j)$ and $j/(i+j)$. These conventions match Mathar [M], Definition 1 and Theorem 2 on page 3, and Definition 4 on page 4. Write $\mathcal{D}_c = \{j : i/j \in \mathcal{F}_c\}$ and $D(c) = |\mathcal{D}_c|$.

Lemma 4.1 (domain). Every node at a positive level is a reduced fraction in $(0,1)$.

Proof. This holds for $1/2$. If $0 < i < j$ and $\gcd(i,j) = 1$, both new numerators are positive and smaller than $i+j$, while

$$\gcd(i, i+j) = \gcd(i,j) = 1, \qquad \gcd(j, i+j) = \gcd(i,j) = 1.$$

Induction proves the assertion. □

Lemma 4.2 (bridge). For $c \geq 1$,

$$\Phi(\mathcal{F}_c) = T_{c-1}, \qquad \mathcal{D}_c = W_c, \qquad D(c) = w_c, \qquad \Phi(i/j) = (i, j-i).$$

Proof. The inverse on the stated domains is $(p,q) \mapsto p/(p+q)$. Let $S(p,q) = (q,p)$ and $A_c = \Phi(\mathcal{F}_c)$. Then $A_1 = \{(1,1)\}$. At every later level the two children of any parent are

complementary, so their images are swapped pairs; hence $S(A_c) = A_c$. Direct substitution gives images $L(p, q)$ and $SL(p, q)$ for the two children. Since $SL = RS$,

$$A_{c+1} = L(A_c) \cup SL(A_c) = L(A_c) \cup R(A_c).$$

Induction gives the level identity. The denominator of i/j is the sum of the coordinates of $\Phi(i/j)$, proving the set and cardinality statements. This is not a pointwise identification of both labelled edges. □

The bridge lemma is used for attribution and comparison, not as a priority claim. Our primary proofs of the spectrum bounds do not require published asymptotics for $D(c)$.

5. Growth bounds and sparsity

Put $\varphi = (1 + \sqrt{5})/2$.

Theorem 5.1 (explicit bound). For all $k \geq 0$,

$$a_k \leq 1 + A(5/3)^k \leq 8(5/3)^k, \qquad A = \frac{95367431640625}{14169773449216} < 7.$$

Proof. Put $r = 5/3$. Assign weights $1, 2, 3, 4, 5, 5$ to primes $2, 3, 5, 7, 11, 13$, extending additively over prime factors. For integers 2 through 13 the weights are respectively

$$1, 2, 2, 3, 3, 4, 3, 4, 4, 5, 4, 5.$$

The least c with $n \leq F_{c+2}$, in the same order, is

$$1, 2, 3, 3, 4, 4, 4, 5, 5, 5, 5, 5.$$

Thus every block of cost $c \leq 5$ has weight at most c. For $c \geq 6$, $F_{c+2} \leq r^c$: bases are $21 \cdot 3^6 = 15309 \leq 15625 = 5^6$ and $34 \cdot 3^7 = 74358 \leq 78125 = 5^7$. The recurrence step uses $1 + r \leq r^2$.

Choose a decomposition (PR) of a positive value. Resolve short blocks into the six primes, and multiply all long blocks into q. If e_p are the six exponents and $E = e_2 + 2e_3 + 3e_5 + 4e_7 + 5e_{11} + 5e_{13}$, then $E \leq k$ and $q \leq r^{k-E}$. For each exponent vector there are at most $\lfloor r^{k-E} \rfloor$ possible positive integers q. Counting descriptions overcounts rather than undercounts values:

$$a_k - 1 \leq \sum_{e:E \leq k} \lfloor r^{k-E} \rfloor \leq r^k \prod_{d \in (1,2,3,4,5,5)} (1 - r^{-d})^{-1} = Ar^k.$$

The sum is over vectors, not distinct costs. No coprimality between q and the six primes is assumed. The displayed rational value follows by multiplying the six rational factors. Zero contributes the separate 1. □

Theorem 5.2 (upper exponent). For every $\lambda > \varphi$ there is a finite C_λ such that $a_k \leq 1 + C_\lambda \lambda^k$ for all $k \geq 0$. Consequently $\limsup a_k^{1/k} \leq \varphi$.

Proof. Induction gives $F_{c+2} \leq \varphi^{c+1}$. Choose $L \geq 2$ with $\varphi(\varphi/\lambda)^L \leq 1$. Blocks of cost $c \geq L$ have size at most λ^c. All shorter blocks lie in $\{2, \ldots, F_{L+1}\}$. Give each short factor artificial cost 1, which cannot exceed its actual cost. Counting their exponent vectors and the remaining positive integer q, as above, gives

$$C_\lambda = (1 - \lambda^{-1})^{-(F_{L+1}-1)}.$$

Taking k-th roots and then $\lambda \downarrow \varphi$ proves the assertion. The constants are not uniform as $\lambda \downarrow \varphi$, so this does not prove $O(\varphi^k)$. □

Corollary 5.3 (sparsity). Define, for $k \geq 1$,

$$M_k = \#\{m \in \mathbb{Z} : 1 \leq m \leq 2^{k-1},\ m \notin V_k\}.$$

Then

$$0 \leq 1 - \frac{M_k}{2^{k-1}} \leq 2A(5/6)^k < 14(5/6)^k, \qquad \frac{M_k}{2^{k-1}} \longrightarrow 1, \qquad \frac{a_k}{2^k} \longrightarrow 0.$$

Proof. The first difference is exactly $|V_k \cap [1, 2^{k-1}]|/2^{k-1}$, at most $(a_k - 1)/2^{k-1}$. Apply Theorem 5.1. Also $a_k/2^k \leq 2^{-k} + A(5/6)^k$. □

The looser displayed rate $14(5/6)^k$ becomes smaller than 1 only for $k \geq 15$; for smaller k, it is valid but does not improve the elementary probability bound 1. This is an asymptotic rate. No formula expressing M_k exactly in terms of a_k and a fixed number of high values is required.

6. A common upper growth exponent

For $n \geq 2$, define $\gamma(n) = \min\{c : n \in W_c\}$. It is finite, since $(n-1, 1) \in T_{n-2}$. Let $\kappa(1) = 0$ and let $\kappa(v)$, for $v \geq 2$, be the minimum of $\sum_i \gamma(g_i)$ over products $v = \prod_i g_i$, $g_i \geq 2$. Minima exist, since there is a one-factor decomposition. From (PR),

$$V_k \setminus \{0\} = \{v \geq 1 : \kappa(v) \leq k\}. \tag{COST}$$

Let $B_c = \{n \geq 2 : \gamma(n) = c\}$, $N_c = |B_c|$, and let $\mathcal{A} = \{n \geq 2 : \kappa(n) = \gamma(n)\}$ be the cost atoms. Set $\alpha_c = |\mathcal{A} \cap B_c|$. These atoms need not be primes. All these layers are finite by Lemma 3.2.

Lemma 6.1 (atomic transfer). If $\lambda > 1$ and $\sum_c \alpha_c \lambda^{-c} < \infty$, then $a_k = O(\lambda^k)$.

Proof. Every optimal factorization uses atoms: otherwise replace a non-atom factor by a strictly cheaper factorization. Put $x = \lambda^{-1}$. The product

$$\mathcal{P}(x) = \prod_{n \in \mathcal{A}} (1 - x^{\gamma(n)})^{-1}$$

is finite, since

$$\log \mathcal{P}(x) \leq (1-x)^{-1} \sum_c \alpha_c x^c.$$

This follows first for finite products from $-\log(1-t) \leq t/(1-x)$ when $0 \leq t \leq x < 1$, and then by monotone passage to the limit. Choose one optimal multiset per value. Distinct values give distinct chosen multisets because their products differ. Descriptions of cost at most k have weight at least x^k, and their total weight is at most $\mathcal{P}(x)$. Hence $a_k - 1 \leq x^{-k}\mathcal{P}(x)$. □

Theorem 6.2 (E).

$$\limsup_{k\to\infty} a_k^{1/k} = \limsup_{c\to\infty} \alpha_c^{1/c} = \limsup_{c\to\infty} N_c^{1/c} = \limsup_{c\to\infty} w_c^{1/c} = \rho, \qquad \frac{2}{\varphi} \leq \rho \leq \varphi.$$

Proof. Write r_a, r_α, r_N, r_w for these four upper rates. Lemma 6.1 and the root test give $r_a \leq \max(1, r_\alpha)$. The inclusions defining the layers give $r_\alpha \leq r_N \leq r_w$. Since $W_c \subseteq V_c$ by (PR), $r_w \leq r_a$.

At depth $c-1$ there are 2^{c-1} distinct ordered pairs. Each sum is at most F_{c+2}, and a fixed sum s has at most $s - 1 \leq F_{c+2}$ positive ordered decompositions. Therefore

$$w_c \geq \frac{2^{c-1}}{F_{c+2}} \geq \frac{(2/\varphi)^c}{2\varphi}.$$

Thus $r_w \geq 2/\varphi > 1$, which removes the threshold 1 from the first inequality and forces equality throughout. Theorem 5.2 supplies the upper bound φ. □

By Lemma 4.2, $\rho = \limsup D(c)^{1/c}$. This substitution alone does not determine ρ. A count of all fractions or the growth of the largest denominator is not the count of distinct denominators.

7. Costs and failures of four halving properties

The cost is subadditive by concatenating factorizations. Since $F_{c+2} \le 2^c$, we also have

$$v \le 2^{\kappa(v)}, \qquad \kappa(2v) \le \kappa(v) + 1.$$

For $k \ge 1$, define R1(k) as $2v \in V_k \Rightarrow v \in V_{k-1}$ for every integer $v \ge 1$. Define R2(k) as $e_k = 0$, where $e_k = |2(V_k \setminus V_{k-1}) \cap V_k|$. Define HP($k$) as: $(a, b) \in K_k$ and $a + b$ even imply $(a+b)/2 \in V_k$. Define HH(j) as: $(a, b) \in K_j$ and b even imply $a+b/2 \in V_{j+1}$. Global versions quantify over all allowed levels.

Lemma 7.1 (R). R1(k) fails at v exactly when $\kappa(2v) \le k \le \kappa(v)$. Consequently global R1 is equivalent to $\kappa(2v) = \kappa(v) + 1$ for every $v \ge 1$.

Proof. The local statement follows from (COST), including integer strictness for $v \notin V_{k-1}$. Equality of costs with increment 1 excludes every such interval. If equality fails, subadditivity and integrality give $\kappa(2v) \le \kappa(v)$. For $v \ge 2$, choose $k = \kappa(v) \ge 1$. For $v = 1$, the costs are 0 and 1, so no exception occurs. □

Proposition 7.2 (finite certificate W-3847). Subject to the exact finite certificate described below, R1(19), R2(19), HP(18), and HH(17) are false.

Proof and finite evidence. Let $S(p, q)$ be the sum of Euclidean quotients for a positive coprime pair. The unique subtractive path to $(1, 1)$ has length $S(p, q) - 1$: the final Euclidean subtraction to zero is excluded from the tree. Thus

$$\gamma(n) = \min_{\substack{1 \le p < n \\ \gcd(p,n)=1}} S(p, n - p).$$

The supplied certificate exhausts all 3846 splits for 3847, reporting minimum 19 and 116 ordered minimizers, including $(672, 3175)$. Trial division by all primes at most 61 proves primality, since $\lfloor \sqrt{3847} \rfloor = 62$. A prime has only its single-factor block decomposition, so $\kappa(3847) = \gamma(3847) = 19$. In particular $3847 \notin V_{18}$ and $3847 \in V_{19}$.

The word $LRLLRLRLRLRRLRRLRL$, of length 18, reaches $(2985, 4709)$, whose sum is $7694 = 2 \cdot 3847$. Its predecessor is $(2985, 1724)$. The complete path is reproduced in Appendix B, which also specifies a complete exact reproduction algorithm; the associated source certificate is preserved in the HYDRA research archive. Hence $7694 \in W_{19} \subseteq V_{19}$. This proves failure of R1(19); membership of 3847 in $V_{19} \setminus V_{18}$ gives $e_{19} \ge 1$, disproving R2(19). Genealogy puts the endpoint pair in K_{18}, disproving HP(18). The predecessor lies in K_{17}, has even second coordinate, and $2985 + 1724/2 = 3847 \notin V_{18}$, disproving HH(17). □

The path is a direct membership proof. Nonmembership depends on an exhaustive finite minimum, not on sampling. The manuscript editor did not repeat that computation, while the later Track B review recomputed the finite minimum and minimizer count with a separate implementation. This is exact finite evidence within its specified verification scope, not an independent certification of the entire article. The assertions do not classify all conceivable halving identities.

8. Discussion and open problems

The product representation links a classical coprime-pair tree to a sparse value spectrum of Pell polynomials. The counting proofs tolerate nonunique factorizations. The exponent identity likewise counts descriptions only as upper bounds; its product is not an exact Euler product for values.

The ordinary limit of $a_k^{1/k}$, equality $\rho = \varphi$, and $a_k = O(\varphi^k)$ remain unsettled here. We refer to these separate research questions as G3, G1, and G2, respectively. If $\rho < \varphi$, atomic

transfer with an intermediate λ gives G2 and negates G1. If $\rho = \varphi$, G1 holds but G2 and G3 are not settled.

For $\lambda > 1$, the inclusions $B_c \subseteq W_c \subseteq \bigcup_{j \le c} B_j$ give

$$\sum_c N_c \lambda^{-c} \le \sum_c w_c \lambda^{-c} \le (1 - \lambda^{-1})^{-1} \sum_c N_c \lambda^{-c}.$$

Thus $\rho < \lambda$ implies convergence, and convergence implies $\rho \le \lambda$. Boundary convergence at $\lambda = \rho$ is not determined by the root test.

Our ten-query supplementary literature search, together with the identified sources, did not locate a proved asymptotic for distinct Kepler denominators that settles these targets. The A294444 page was inaccessible in the retrieval tool; Mathar supplies its identification and finite table. Results about random paths, sums over nodes, or normality of continued fractions do not by themselves settle this distinct-value question. This is a limitation of search, not a theorem of absence.

9. References

[M] Richard J. Mathar, The Kepler Binary Tree of Reduced Fractions, 2017. `https://oeis.org/A294443/a294443.pdf`. Definitions and level convention checked on pp. 3–4; denominator counts in Section 3.

[O] OEIS Foundation Inc., A086592, Denominators in left-hand half of Kepler's tree of fractions. `https://oeis.org/A086592`. Accessed 6 November 2026; levelwise comment by Yosu Yurramendi, 6 April 2016.

[D] Karma Dajani, Mathijs R. de Lepper, and E. Arthur Robinson Jr., Introducing Minkowski normality, Journal of Number Theory 211 (2020), 455–476; arXiv:1902.10373. `https://arxiv.org/abs/1902.10373`. Cited for related Kepler/continued-fraction context, not a distinct-denominator asymptotic.

[P] Extremal coefficients of Pell representation polynomials and integer fibres in the silver ratio, aiXiv:261101.000005, version 1.0, 2026. `https://aixiv.science/abs/aixiv.261101.000005`. Page credits openai/gpt-6-luna and anthropic/claude-opus-5.5, with corresponding human Marcin Zwierzyński; only page and abstract were inspected in the source review.

Appendix A. Further structural conventions and errata

Define $b_j = |\{v \ge 1 : \kappa(v) = j\}|$, so $b_0 = 1$, $a_k = 1 + \sum_{j \le k} b_j$. For the formal series $\mathcal{A}(z) = \sum a_k z^k$, $\mathcal{B}(z) = \sum b_k z^k$,

$$\mathcal{A}(z) = \frac{1 + \mathcal{B}(z)}{1 - z}.$$

The two sequences have the same limsup exponential rate: $b_k \le a_k$ gives one inequality, while bounding b_j by $C\lambda^j$ for any $\lambda > \max(1, \limsup b_j^{1/j})$ gives the reverse after summation. Theorem E ensures the rate exceeds 1. Their common radius is $1/\rho < 1$. No equality of liminf is established.

The cost κ is not additive on coprime factors: the historical finite data give costs $7, 9, 15$ for $23, 53, 1219 = 23 \cdot 53$. This concerns a particular prime-cost product, not the impossibility of all formal Euler products. That optional finite observation is not needed for the main theorems.

For completeness define the continuant $K(\varnothing) = 1$, $K(a_1) = a_1$, with recurrence $K(a_1, \ldots, a_r) = a_r K(a_1, \ldots, a_{r-1}) + K(a_1, \ldots, a_{r-2})$. Here this K is a continuant,

not the indexed pair set K_k. Standard fraction multiplication gives $[a_1;\ldots;a_r] = K(a_1,\ldots,a_r)/K(a_2,\ldots,a_r)$. Then

$$W_c = \{K(a_1,\ldots,a_r) + K(a_2,\ldots,a_r) : a_i \geq 1,\ \sum_i a_i = c\}.$$

Indeed swap coordinates to make $p \geq q$, preserving the tree and the sum. The Euclidean quotient sum is one more than the subtractive depth. A canonical final quotient at least 2 can be replaced by the pair consisting of one less and 1, preserving both fraction and quotient sum. Conversely a terminal 1 can be merged with its predecessor. The exceptional fraction 1 is [1], giving $W_1 = \{2\}$. This proves both inclusions. The number of quotients is not word length; $[c]$ has one quotient but depth $c-1$.

The mandatory historical corrections are respected as follows. M-1 replaces old OPEN statuses for the four global halving properties by the four falsifications of Section 7. M-2 uses only the exact lower bound in Theorem E. M-3 excludes the unproved liminf assertion. M-4 separates G1 from G2 on the boundary. M-5 treats an odd-block liminf criterion only as sufficient, never necessary. M-6 leaves boundary summability open. No numerical fit is a global asymptotic theorem.

Appendix B. Finite verification and disclosure

A minimal independent verification of Proposition 7.2 can use the following exact algorithm. Test every integer divisor from 2 to 62 for 3847. For each $p = 1,\ldots,3846$, evaluate $S(p, 3847-p)$ by repeated integer division: add the quotient of the larger coordinate by the smaller and replace the larger by the remainder, until one coordinate vanishes. Retain all minimizers. Verify that the minimum is 19 and that there are 116 ordered minimizers. Finally apply the 18-letter word from Section 7 using integer addition. This completely specifies the required finite test and its scope. The editorial phase did not re-execute it; the subsequent Track B review documented a separate exact computation.

The path is

$$(1,1),(1,2),(3,2),(3,5),(3,8),(11,8),(11,19),$$
$$(30,19),(30,49),(79,49),(79,128),(207,128),$$
$$(335,128),(335,463),(798,463),(1261,463),$$
$$(1261,1724),(2985,1724),(2985,4709).$$

The source certificate SHA-256 is `2AF1C46BF4F1EF67FC0E9073E541E8E45D38AB1446774AAF90923B5D697800E5`.

The mathematical arguments originate in the HYDRA C1/C5/C6 research and adjudication records, with the governing D5/D6 qualifications retained in the research provenance archive. C1 involved GPT-6; C5 used Codex with its historical model version unresolved, and C6 used DeepSeek Harness with its exact model unresolved. Claude Opus 5.5 contributed to adjudications and to the exponent, cost, and witness arguments. GPT-6.1 Sol prepared the original R3 editorial edition in Codex desktop. Kimi Work reviewed R3 in Track B; it had previously authored C3 and therefore did not establish independence for the foundational claims. B-2 used separately tasked GPT subagents (recorded as GPT-6.1 Terra, GPT-6 Astra, and GPT-6.1 Sol), but shared filesystem/tool access prevented an enforced blind-phase certification; model-family independence was also not independently established. These audits identified no material proof defect within the reviewed scope but are not human peer review or machine-checked formal verification. The precise identity of some historical runs and parts of the human contribution record remain unknown. Formal authorship, correspondence, and licence terms are specified in separate Owner-approved release metadata, and no priority claim is made.